Online Appendix for "Justified Communication Equilibrium"

Daniel Clark and Drew Fudenberg

OA.1 Equivalent Definition of JCE

We show that it would be equivalent to define JCE by setting $\Theta^{\dagger}(s,\pi) = \{\theta \in \Theta : \widetilde{D}^0_\theta(s,\pi) \not\subseteq \cup_{\theta' \neq \theta} \widetilde{D}_{\theta'}(s,\pi)\}$, rather than $\Theta^{\dagger}(s,\pi) = \{\theta \in \Theta : \widetilde{D}_\theta(s,\pi) \cup \widetilde{D}^0_\theta(s,\pi) \not\subseteq \cup_{\theta' \neq \theta} \widetilde{D}_{\theta'}(s,\pi)\}$.

For every $s \in S$ and $\pi \in \Pi_1 \times \Pi_2$, let

$$\Theta^{\dagger\prime}(s,\pi) = \{\theta \in \Theta : \widetilde{D}^0_\theta(s,\pi) \not\subseteq \cup_{\theta' \neq \theta} \widetilde{D}_{\theta'}(s,\pi)\}$$

be the set of types θ where there is some mixed receiver action $\alpha \in \Delta(BR(\Theta, s))$ that makes θ indifferent between (s,α) and their outcome under π and makes no other type θ' strictly prefer (s,α) to their outcome under π. Additionally, let

$$\overline{\Theta}'(s,\pi) = \begin{cases} \Theta^{\dagger\prime}(s,\pi) & \text{if } \Theta^{\dagger\prime}(s,\pi) \neq \emptyset \\ \Theta & \text{if } \Theta^{\dagger\prime}(s,\pi) = \emptyset \end{cases}.$$

Proposition OA 1. *If π is a PBE-H, then $\overline{\Theta}(s,\pi) = \overline{\Theta}'(s,\pi)$ for all $s \in S$.*

Proof. Fix PBE-H π. We will argue that $\Theta^{\dagger'}(s,\pi) = \Theta^{\dagger}(s,\pi)$, which gives $\overline{\Theta}'(s,\pi) = \overline{\Theta}(s,\pi)$.

First, suppose that $\theta \in \Theta^{\dagger'}(s,\pi)$. Then by definition, $\widetilde{D}^0_\theta(s,\pi) \not\subseteq \cup_{\theta'\neq\theta}\widetilde{D}_{\theta'}(s,\pi)$. Hence, $\widetilde{D}_\theta(s,\pi) \cup \widetilde{D}^0_\theta(s,\pi) \not\subseteq \cup_{\theta'\neq\theta}\widetilde{D}_{\theta'}(s,\pi)$, so $\theta \in \Theta^{\dagger}(s,\pi)$.

Now, suppose that $\theta \in \Theta^{\dagger}(s,\pi)$. Then by definition, $\widetilde{D}_\theta(s,\pi)\cup\widetilde{D}^0_\theta(s,\pi) \not\subseteq \cup_{\theta'\neq\theta}\widetilde{D}_{\theta'}(s,\pi)$. Thus, there is some $\alpha \in \Delta(BR(\Theta,s))$ such that $u_1(\theta,s,\alpha) \geq u_1(\theta,\pi)$ and $u_1(\theta,s,\alpha) \leq u_1(\theta',\pi)$ for all $\theta' \neq \theta$. Since π is a PBE-H, there is also some $\alpha' \in \Delta(BR(\Theta,s))$ such that $u_1(\theta',s,\alpha') \leq u_1(\theta,\pi)$ for all $\theta' \in \Theta$. By continuity, there is some $\nu \in [0,1]$ and $\alpha'' = \nu\alpha + (1-\nu)\alpha'$ such that $u_1(\theta,s,\alpha'') = u_1(\theta,\pi)$, while $u_1(\theta',s,\alpha'') \leq u_1(\theta',\pi)$ for all $\theta' \neq \theta$. As $\alpha'' \in \Delta(BR(\Theta,s))$, it follows that $\widetilde{D}^0_\theta(s,\pi) \not\subseteq \cup_{\theta'\neq\theta}\widetilde{D}_{\theta'}(s,\pi)$, so $\theta \in \Theta^{\dagger'}(s,\pi)$. ■

OA.2 Omitted Analysis of Learning Model

OA.2.1 Continuity of Aggregate Response Mapping

We begin by formally defining the auxiliary maps $\mathscr{L}_1^{\delta,\gamma_1} : \Pi_2 \to (\Delta(\mathcal{H}_1))^\Theta$ and $\mathscr{L}_2^{\gamma_2} : \Pi_1 \to \Delta(\mathcal{H}_2)$ introduced in Appendix D. For each $\theta \in \Theta$, let

$$\mathscr{L}_\theta^{\delta,\gamma_1}(\pi_2)[\emptyset] = 1-\gamma_1,$$
$$\mathscr{L}_\theta^{\delta,\gamma_1}(\pi_2)[(h_1,(s,m,a))] = \gamma_1\mathscr{L}_\theta^{\delta,\gamma_1}(\pi_2)[h_1]i_\theta^{\delta,\gamma_1}(h_1,s,m)\pi_2[a|s,m],$$

for all $h_1 \in \mathcal{H}_1$, $s \in S$, $m \in M$, and $a \in A$. To define $\mathscr{L}_2^{\gamma_2}$, let

$$\mathscr{L}_2^{\gamma_2}(\pi_1)[\emptyset] = 1-\gamma_2,$$
$$\mathscr{L}_2^{\gamma_2}(\pi_1)[(h_2,(\theta,s,m))] = \gamma_2\mathscr{L}_2^{\gamma_2}(\pi_1)[h_2]\lambda(\theta)\pi_1[s,m|\theta],$$

for all $h_2 \in \mathcal{H}_2$, $\theta \in \Theta$, $s \in S$, and $m \in M$.

We now establish the continuity of various mappings involving distributions over histories, which we endow with the sup-norm topology.

Claim OA 1. *The aggregate strategy mapping* $\sigma^{\delta,\gamma_1} : (\Delta(\mathcal{H}_1))^{\Theta} \times \Delta(\mathcal{H}_2) \to \Pi_1 \times \Pi_2$ *is continuous.*

Proof. We prove that $\sigma_1^{\delta,\gamma_1} : (\Delta(\mathcal{H}_1))^{\Theta} \to \Pi_1$ is continuous. An analogous argument handles $\sigma_2 : \Delta(\mathcal{H}_2) \to \Pi_2$.

To show that $\sigma_1^{\delta,\gamma_1}$ is continuous, we establish that $\lim_{\mu_1' \to \mu_1} \sigma_1^{\delta,\gamma_1}(\mu_1')[s, m|\theta] = \sigma_1^{\delta,\gamma_1}(\mu_1)[s, m|\theta]$ for all $s \in S$, $m \in M$, $\theta \in \Theta$, and $\mu_1 \in (\Delta(\mathcal{H}_1))^{\Theta}$. Since $\sum_{s,m} \sigma_1^{\delta,\gamma_1}(\mu_1')[s, m|\theta] = 1$ for all $\mu_1 \in (\Delta(\mathcal{H}_1))^{\Theta}$, it suffices to show that $\liminf_{\mu_1' \to \mu_1} \sigma_1^{\delta,\gamma_1}(\mu_1')[s, m|\theta] \geq \sigma_1^{\delta,\gamma_1}(\mu_1)[s, m|\theta]$ for all s, m, and θ. For any $\varepsilon > 0$, let $\mathcal{H}_{1,\varepsilon}$ be a finite set of sender histories such that $\sum_{h_1 \in \mathcal{H}_{1,\varepsilon} : \mathbf{x}_{\theta}^{\delta,\gamma_1}(h_1) = (s,m)} \mu_{\theta}[h_1] \geq \sigma_1^{\delta,\gamma_1}(\mu_1)[s, m|\theta] - \varepsilon$. By the nature of the sup-norm topology, $\lim_{\mu_1' \to \mu_1} \sum_{h_1 \in \mathcal{H}_{1,\varepsilon} : \mathbf{x}_{\theta}^{\delta,\gamma_1}(h_1) = (s,m)} \mu_{\theta}'[h_1] = \sum_{h_1 \in \mathcal{H}_{1,\varepsilon} : \mathbf{x}_{\theta}^{\delta,\gamma_1}(h_1) = (s,m)} \mu_{\theta}[h_1]$. Since $\mu_{\theta}'[h_1] \geq 0$ for all $h_1 \in \mathcal{H}_1$ and $\mu_1' \in (\Delta(\mathcal{H}_1))^{\Theta}$, it follows that $\liminf_{\mu_1' \to \mu_1} \sigma_1^{\delta,\gamma_1}(\mu_1')[s, m|\theta] = \liminf_{\mu_1' \to \mu_1} \sum_{h_1 : \mathbf{x}_{\theta}^{\delta,\gamma_1}(h_1) = (s,m)} \mu_{\theta}'[h_1] \geq \lim_{\mu_1' \to \mu_1} \sum_{h_1 \in \mathcal{H}_{1,\varepsilon} : \mathbf{x}_{\theta}^{\delta,\gamma_1}(h_1) = (s,m)} \mu_{\theta}'[h_1] \geq \sigma_1^{\delta,\gamma_1}(\mu_1)[s, m|\theta] - \varepsilon$. As this holds for arbitrary $\varepsilon > 0$, the desired conclusion follows. ■

Claim OA 2. *Both* $\mathscr{L}_1^{\delta,\gamma_1} : \Pi_2 \to (\Delta(\mathcal{H}_1))^{\Theta}$ *and* $\mathscr{L}_2^{\gamma_2} : \Pi_1 \to \Delta(\mathcal{H}_2)$ *are continuous.*

Proof. We prove that $\mathscr{L}_1^{\delta,\gamma_1} : \Pi_2 \to (\Delta(\mathcal{H}_1))^{\Theta}$ is continuous. An analogous argument handles $\mathscr{L}_2^{\gamma_2} : \Pi_1 \to \Delta(\mathcal{H}_2)$.

For all $\pi_2 \in \Pi_2$, $\mathscr{L}_1^{\delta,\gamma_1}(\pi_2)[h_1] \leq (1 - \gamma_1)\gamma_1^t$ for every history h_1 of length t. Since $\lim_{t \to \infty}(1 - \gamma_1)\gamma_1^t = 0$, to establish that $\mathscr{L}_1^{\delta,\gamma_1}(\pi_2)$ is a continuous function of π_2, it thus suffices to show that $\mathscr{L}_1^{\delta,\gamma_1}(\pi_2)[h_1]$ is continuous for every history $h_1 \in \mathcal{H}_1$. We show this inductively over sender histories. For the null sender history $h_1 = \emptyset$, $\mathscr{L}_1^{\delta,\gamma_1}(\pi_2)[\emptyset]$ for all $\pi_2 \in \Pi_2$ and is thus continuous. Assuming that $\mathscr{L}_1^{\delta,\gamma_1}(\pi_2)[h_1]$ is a continuous function of π_2, it follows that $\mathscr{L}_1^{\delta,\gamma_1}(\pi_2)[(h_1, (s, m, a))]$ is a continuous function of π_2 for all s, m, and a, as can be seen from the expression for $\mathscr{L}_1^{\delta,\gamma_1}$ given earlier. This completes the inductive argument. ■

Corollary OA 1. *The aggregate response mapping* $\mathscr{R}^{\delta,\gamma_1,\gamma_2} : \Pi_1 \times \Pi_2 \to \Pi_1 \times \Pi_2$ *is continuous.*

Proof. By Claims OA 1 and OA 2, $\sigma_1^{\delta,\gamma_1}$ and $\mathscr{L}_1^{\delta,\gamma_1}$ are continuous. Thus $\mathscr{R}_1^{\delta,\gamma_1}(\pi_2) = \sigma_1^{\delta,\gamma_1}(\mathscr{L}_1^{\delta,\gamma_1}(\pi_2))$ is a continuous function of π_2. Likewise, since σ_2 and $\mathscr{L}_2^{\gamma_2}$ are continuous, $\mathscr{R}_2^{\gamma_2}(\pi_1) = \sigma_2(\mathscr{L}_2^{\gamma_2}(\pi_1))$ is a continuous function of π_1. ■

OA.2.2 Characterization of Steady State Profiles

Proposition OA 2. *Strategy profile π is a fixed point of $\mathscr{R}^{\delta,\gamma_1,\gamma_2}$ if and only if there is some steady state μ such that $\sigma^{\delta,\gamma_1}(\mu) = \pi$.*

Proof. Suppose that μ is a steady state satisfying $\sigma^{\delta,\gamma_1}(\mu) = \pi$. Since μ is a steady state, the aggregate receiver play in every period is fixed at $\pi_2 = \sigma_2(\mu)$. By definition, $\mathscr{L}_1^{\delta,\gamma_1}(\pi_2)$ is the $t \to \infty$ limit of the distribution over histories in the sender population when the aggregate receiver play is fixed at π_2. Since μ is a steady state, it follows that $\mathscr{L}_1^{\delta,\gamma_1}(\pi_2) = \mu_1$. From this, we obtain $\mathscr{R}_1^{\delta,\gamma_1}(\pi_2) = \sigma_1^{\delta,\gamma_1}(\mathscr{L}_1^{\delta,\gamma_1}(\pi_2)) = \sigma_1^{\delta,\gamma_1}(\mu_1) = \pi_1$. An almost identical argument shows that $\mathscr{R}_2^{\gamma_2}(\pi_1) = \pi_2$. We conclude that $\mathscr{R}^{\delta,\gamma_1,\gamma_2}(\pi) = \pi$.

Conversely,suppose that π is a fixed point of $\mathscr{R}^{\delta,\gamma_1,\gamma_2}$. Let μ be the state given by $\mu_1 = \mathscr{L}_1^{\delta,\gamma_1}(\pi_2)$ and $\mu_2 = \mathscr{L}_2^{\gamma_2}(\pi_1)$. Observe that $\sigma_1^{\delta,\gamma_1}(\mu_1) = \sigma_1^{\delta,\gamma_1}(\mathscr{L}_1^{\delta,\gamma_1}(\pi_2)) = \mathscr{R}_1^{\delta,\gamma_1}(\pi_2) = \pi_1$ and $\sigma_2(\mu_2) = \sigma_2(\mathscr{L}_2^{\gamma_2}(\pi_1)) = \mathscr{R}_2^{\gamma_2}(\pi_1) = \pi_2$, so $\pi = \sigma^{\delta,\gamma_1}(\mu)$ is the aggregate strategy profile for state μ. All that remains is to establish that μ is a steady state, which amounts to showing that $\mathbf{f}_\theta^{\delta,\gamma_1}(\mathscr{L}_1^{\delta,\gamma_1}(\pi_2))[h_1] = \mathscr{L}_1^{\delta,\gamma_1}(\pi_2)[h_1]$ for all $h_1 \in \mathcal{H}_1$ and $\theta \in \Theta$ and $\mathbf{f}_2^{\delta,\gamma_1,\gamma_2}(\mathscr{L}_2^{\gamma_2}(\pi_1))[h_2] = \mathscr{L}_2^{\gamma_2}(\pi_1)[h_2]$ for all $h_2 \in \mathcal{H}_2$. We argue inductively over sender histories that $\mathbf{f}_\theta^{\delta,\gamma_1}(\mathscr{L}_1^{\delta,\gamma_1}(\pi_2))[h_1] = \mathscr{L}_1^{\delta,\gamma_1}(\pi_2)[h_1]$ for all $h_1 \in \mathcal{H}_1$. (A similar inductive argument shows that $\mathbf{f}_2^{\delta,\gamma_1,\gamma_2}(\mathscr{L}_2^{\gamma_2}(\pi_1))[h_2] = \mathscr{L}_2^{\gamma_2}(\pi_1)[h_2]$ for all $h_2 \in \mathcal{H}_2$.) For the null sender history $h_1 = \emptyset$, the equality holds since $\mathbf{f}_\theta^{\delta,\gamma_1}(\mathscr{L}_1^{\delta,\gamma_1}(\pi_2))[\emptyset] = 1 - \gamma_1 = \mathscr{L}_1^{\delta,\gamma_1}(\pi_2)[\emptyset]$. Assuming that $\mathbf{f}_\theta^{\delta,\gamma_1}(\mathscr{L}_1^{\delta,\gamma_1}(\pi_2))[h_1] = \mathscr{L}_1^{\delta,\gamma_1}(\pi_2)[h_1]$ holds, it necessarily follows that $\mathbf{f}_\theta^{\delta,\gamma_1}(\mathscr{L}_1^{\delta,\gamma_1}(\pi_2))[(h_1,(s,m,a))] = \mathscr{L}_1^{\delta,\gamma_1}(\pi_2)[(h_1,(s,m,a))]$ for all s, m, and a since $\sigma_2(\mu_2) = \pi_2$. This completes the inductive argument. ■

OA.3 Comparison with RCE

In this section, we restrict attention to signaling games without communication, i.e. M is singleton. We write $\Pi_2^\bullet = \times_{s \in S} \Delta(BR(\Theta, s))$ for the set of receiver strategies that assign probability 0 to conditionally dominated responses.

Definition OA 1 (Fudenberg and He, 2020). *Signal $s \in S$ is* **more rationally-compatible** *with θ' than θ'', written as $\theta' \succsim_s \theta''$,*

$$u_1(\theta'', s, \pi_2(\cdot|s)) \geq \max_{s' \neq s} u_1(\theta'', s', \pi_2(\cdot|s')) \text{ implies that}$$
$$u_1(\theta', s, \pi_2(\cdot|s)) > \max_{s' \neq s} u_1(\theta', s', \pi_2(\cdot|s')).$$

In words, this says that type θ' is more rationally-compatible with signal s than is θ'' if any undominated receiver strategy that makes θ'' willing to play s makes θ' strictly prefer to play it. Let $P_{\theta' \triangleright \theta''} = \{p \in \Delta(\Theta) : \lambda(\theta'')p(\theta') \geq \lambda(\theta')p(\theta'')\}$ be the set of probability distributions over sender type where the odds ratio of θ' to θ'' exceed their odds ratio under the prior distribution. For $s \in S$ and $\pi \in \Pi_1 \times \Pi_2$, let $\overline{P}(s, \pi) \subseteq \Delta(\Theta)$ be the set of beliefs over the sender type given by

$$\overline{P}(s, \pi) = \begin{cases} \Delta(E(s, \pi)) \cap \left(\cap_{(\theta', \theta'') \text{ s.t. } \theta' \succsim_s \theta''} P_{\theta' \triangleright \theta''}\right) & \text{if } E(s, \pi) \neq \emptyset \\ \Delta(\Theta) & \text{if } E(s, \pi) = \emptyset \end{cases},$$

and let $BR(\overline{P}(s, \pi), s) = \cup_{p \in \overline{P}(s, \pi)} BR(p, s)$ be the set of receiver best responses to signal s for some $p \in \overline{P}(s, \pi)$.

Definition OA 2 (Fudenberg and He, 2020). *Strategy profile π is a* **rationality-compatible equilibrium (RCE)** *if it is a PBE-H where, for every $s \in S$, $\pi_2(\cdot|s) \in \Delta(BR(\overline{P}(s, \pi), s))$.*

This definition requires that the receiver's posterior likelihood ratio for types θ' and

θ'' dominates the prior likelihood ratio whenever $\theta' \succsim_s \theta''$. It also requires that the posterior assigns probability 0 to equilibrium-dominated types.

Proposition OA 3. *If π is a justified communication equilibrium, then π is an RCE.*

Intuitively, any response that makes a less compatible type weakly prefer to play s makes more compatible types strictly prefer to play it, so less compatible types are not justified.

Proof. Fix $s \in S$. We will argue that $\Delta(\overline{\Theta}(s,\pi)) \subseteq \overline{P}(s,\pi)$. Thus any $\alpha \in \Delta(BR(\overline{\Theta}(s,\pi),s))$ also belongs to $\Delta(BR(\overline{P}(s,\pi),s))$. Consequently, the justified response criterion of JCE along with the fact that every JCE is a PBE-H implies that π is an RCE.

Since $\Delta(\overline{\Theta}(s,\pi)) \subseteq \Delta(\Theta) = \overline{P}(s,\pi)$ when $E(s,\pi) = \emptyset$, we need only handle the case where $E(s,\pi) \neq \emptyset$. In this case by Lemma A1, $\overline{\Theta}(s,\pi) = \Theta^{\dagger}(s,\pi)$ and $\Delta(\overline{\Theta}(s,\pi)) \subseteq \Delta(E(s,\pi))$. Suppose that θ' and θ'' are two types such that $\theta' \succsim_s \theta''$. Then Definition OA 2 implies that $\widetilde{D}_{\theta''}(s,\pi) \cup \widetilde{D}^0_{\theta''}(s,\pi) \subseteq \widetilde{D}_{\theta'}(s,\pi)$, so $\theta'' \notin \Theta^{\dagger}(s,\pi)$. As a result, $\Delta(\overline{\Theta}(s,\pi)) = \Delta(\Theta^{\dagger}(s,\pi)) \subseteq \cap_{(\theta',\theta'') \text{ s.t. } \theta' \succsim_s \theta''} P_{\theta' \triangleright \theta''}$. We conclude $\Delta(\overline{\Theta}(s,\pi)) \subseteq \Delta(E(s,\pi)) \cap (\cap_{(\theta',\theta'') \text{ s.t. } \theta' \succsim_s \theta''} P_{\theta' \triangleright \theta''}) = \overline{P}(s,\pi)$. ■

OA.4 Proof of Proposition C1

Proposition C1. *If π is a uniformly justified JCE in a strictly monotonic signaling game, it induces the same distribution over $\Theta \times S \times A$ as a stable profile for all non-doctrinaire priors g_1, g_2, including those that do not satisfy initial trust.*

Proof. Because π is a uniformly justified JCE in a strictly monotonic signaling game, $\pi_2(\cdot|s,m) = \pi_2(\cdot|s,m')$ for all $s \in S$ and $m, m' \in M$ such that $(s,m),(s,m') \in X^{\text{on}}$. Thus, for every $s \in S^{\text{on}}$, there is some $a_s \in A$ such that $\pi_2(a_s|s,m) = 1$ for all $(s,m) \in X^{\text{on}}$. For all $s \in S^{\text{off}}$, fix some $a_s \in BR(\overline{\Theta}(s,\pi),s)$.

Our construction modifies the aggregate receiver response so that the response to any s is a_s with high probability unless the aggregate sender play is such that each type $\theta \in \Theta$ uses s_θ with sufficiently high probability. We show that the fixed points of this modified aggregate response mapping correspond to fixed points of the true aggregate response mapping in the iterated limit where $\gamma_1 \to 1$ then $\delta \to 1$ then $\gamma_2 \to 1$. Moreover, we show that the limit of these steady state profiles induce the same distribution over $\Theta \times S \times A$ as π.

Because π is a uniformly justified JCE in a strictly monotonic signaling game, there is an $\varepsilon > 0$ such that the following two properties hold. First, when $\pi_2(a_s|s,m) \geq 1-\varepsilon$ for all s, playing s_θ paired with message m is strictly better for type θ than playing any other $s' \neq s_\theta$ paired with any m'. Second, if $\pi_1(s_\theta, m|\theta) \geq 1-\varepsilon$ for every $\theta \in \Theta$, it is strictly optimal for the receiver to respond to (s,m) with a_s for every $s \in S^{\text{on}}$. Fix such an ε.

Let $\kappa : \mathbb{R} \to [0,1]$ be a continuous function such that $\kappa(z) = 0$ for all $z \leq 0$ and $\kappa(z) = 1$ for all $z \geq 1$. Also, let $\phi : \Pi_1 \times \Pi_2 \to \Pi_2$ be the mapping

$$\phi(\pi_1,\pi_2)(\cdot|s,m) = \left(1 - \kappa\left(\frac{2}{\varepsilon}(\min_{\theta\in\Theta} \pi_1(s_\theta|\theta) - 1 + \varepsilon)\right)\right) \mathbb{1}_{a_s}(\cdot) + \kappa\left(\frac{2}{\varepsilon}(\min_{\theta\in\Theta} \pi_1(s_\theta|\theta) - 1 + \varepsilon)\right) \pi_2(\cdot|s,m)$$

for all $s \in S$ and $m \in M$. Note that ϕ is continuous. Additionally, $\phi(\pi_1,\pi_2)(a_s|s,m) = 1$ when $\pi_1(s_\theta|\theta) \leq 1-\varepsilon$ for some $\theta \in \Theta$, and $\phi(\pi_1,\pi_2) = \pi_2$ when $\pi_1(s_\theta|\theta) \geq 1-\varepsilon/2$ for all $\theta \in \Theta$.

Consider the correspondence $\widetilde{\mathscr{R}}^{\delta,\gamma_1,\gamma_2} : \Pi_1 \times \Pi_2 \to \Pi_1 \times \Pi_2$ given by $\widetilde{\mathscr{R}}^{\delta,\gamma_1,\gamma_2}(\pi) = (\mathscr{R}_1^{\delta,\gamma_1}(\pi_2), \phi(\pi_1, \mathscr{R}_2^{\gamma_2}(\pi_1)))$. Since $\widetilde{\mathscr{R}}^{\delta,\gamma_1,\gamma_2}$ is continuous, Brouwer's fixed point theorem guarantees the existence of a fixed point $(\pi_1^{\delta,\gamma_1,\gamma_2}, \pi_2^{\delta,\gamma_1,\gamma_2})$. We will establish that, in the iterated limit where $\gamma_1 \to 1$ then $\delta \to 1$ then $\gamma_2 \to 1$, $\pi^{\delta,\gamma_1,\gamma_2} = (\pi_1^{\delta,\gamma_1,\gamma_2}, \pi_2^{\delta,\gamma_1,\gamma_2})$ induces the same distribution over $\Theta \times S \times A$ as π. Towards this end, consider a sequence $\{\gamma_{2,j}\}_{j\in\mathbb{N}}$, sequences $\{\delta_{j,k}\}_{j,k\in\mathbb{N}}$, and sequences $\{\gamma_{1,j,k,l}\}_{j,k,l\in\mathbb{N}}$ such that (1) $\lim_{j\to\infty} \gamma_{2,j} = 1$, (2) $\lim_{k\to\infty} \delta_{j,k} = 1$ for all j, (3) $\lim_{l\to\infty} \gamma_{1,j,k,l} = 1$ for all j,k, and (4) $\lim_{j\to\infty} \lim_{k\to\infty} \lim_{l\to\infty} \pi^{\delta_{j,k},\gamma_{1,j,k,l},\gamma_{2,j}} = \pi'$ for some $\pi' = (\pi_1', \pi_2') \in \Pi_1 \times \Pi_2$.

We first establish that $\pi_1'(s_\theta|\theta) \geq 1-\varepsilon$ for all $\theta \in \Theta$. If instead there were some $\theta \in \Theta$ such that $\pi'(s_\theta|\theta) < 1-\varepsilon$, then by construction, $\pi_2'(a_s|s,m) \geq 1-\varepsilon$ for all

$s \in S$ and $m \in M$. Lemma B1 thus requires that $\pi_1'(s_\theta|\theta) = 1$ for all $\theta \in \Theta$, which is a contradiction.

Next we show that $\pi_2'(a_s|s, m) = 1$ for all $s \in S^{\text{on}}$ and $m \in M$ such that $\pi_1'(s, m|\theta) > 0$ for some $\theta \in \Theta$. Fix $s \in S^{\text{on}}$. Consider $m, m' \in M$ such that $\pi_1'(s, m|\theta) > 0$ and $\pi_1'(s, m'|\theta') > 0$ for some $\theta, \theta' \in \Theta$. The construction of $\widetilde{\mathscr{R}}^{\delta,\gamma_1,\gamma_2}$, along with an argument almost identical to the proof of Lemma 2, implies that there exists some $\xi \in [0, 1]$ and $\alpha, \alpha' \in MBR(\Theta, s)$ such that $\pi_2'(\cdot|s, m) = (1 - \xi)\mathbb{1}_{a_s}(\cdot) + \xi\alpha$ and $\pi_2'(\cdot|s, m') = (1 - \xi)\mathbb{1}_{a_s}(\cdot) + \xi\alpha'$. In fact, α and α' must be optimal responses to s under the posterior distributions obtained by updating λ using $\{\pi_1'(s, m|\theta)\}_{\theta\in\Theta}$ and $\{\pi_1'(s, m'|\theta)\}_{\theta\in\Theta}$, respectively. Because the game is strictly monotonic, Lemma B1 implies that $\alpha = \alpha'$. Thus, for a given s, $\pi_2'(\cdot|s, m)$ is the same for all $m \in M$ for which there is a $\theta' \in \Theta$ such that $\pi_1'(s, m|\theta') > 0$. Combining this with the fact that $\pi_1'(s_\theta|\theta) \geq 1 - \varepsilon$ for all θ, it follows that $\pi_2'(a_s|s, m) = 1$ for all $m \in M$ such that $\pi_1'(s, m|\theta) > 0$ for some $\theta \in \Theta$.

Since $\pi_2'(a_s|s, m) = 1$ for all $s \in S^{\text{on}}$ and $m \in M$ such that $\pi_1'(s, m|\theta) > 0$ for some $\theta \in \Theta$, it follows from Lemma B1 that $\pi_1'(s|\theta) = 0$ whenever $s \in S^{\text{on}}$ and $s \neq s_\theta$. We now show that for all $\theta \in \Theta$, $\pi_1'(s|\theta) = 0$ for all $s \in S^{\text{off}}$. Note that, because $\pi_1(s_\theta|\theta) > 0$ for all $\theta \in \Theta$ and $\pi_2(a_{s_\theta}|s_\theta, m) = 1$ for all $\theta \in \Theta$ and $m \in M$ where $\pi_1(s_\theta, m|\theta) > 0$, Lemma B1 implies that $u_1(\theta, \pi') = u_1(\theta, s_\theta, a_{s_\theta}) = u_1(\theta, \pi)$ for all $\theta \in \theta$. Additionally, Lemma B1 requires that $u_1(\theta, s, \pi_2'(\cdot|s, m)) \leq u_1(\theta, \pi') = u_1(\theta, \pi)$ for all $\theta \in \Theta$, $s \in S$, and $m \in M$. Now, suppose that there is some $s \in S^{\text{off}}$ and $m \in M$ such that $\pi_1'(s, m|\theta) > 0$ for some $\theta \in \Theta$. There are two possible cases: (1) There is some $\theta \notin \overline{\Theta}(s, \pi)$ such that $\pi_1'(s, m|\theta) > 0$, and (2) All θ with $\pi_1'(s, m|\theta) > 0$ belong to $\overline{\Theta}(s, \pi)$. In Case (1), because $\pi_2'(\cdot|s, m) \in \Delta(BR(\Theta, s))$, there must be some $\theta' \in \overline{\Theta}(s, \pi)$ such that $u_1(\theta', s, \pi_2'(\cdot|s, m)) > u_1(\theta', \pi)$, which is a contradiction. In Case (2), the construction of $\mathscr{R}^{\delta,\gamma_1,\gamma_2}$, combined with an almost identical argument to the one behind Lemma 2, implies that $\pi_2'(\cdot|s, m) \in \Delta(BR(\overline{\Theta}(s, \pi), s))$. Since π is a uniformly justified JCE, it follows that $u_1(\theta, s, \pi_2'(\cdot|s, m)) < u_1(\theta, \pi)$ for all $\theta \in \Theta$, but this, along with Lemma B1, implies that $\pi_1'(s, m|\theta) = 0$ for all $\theta \in \Theta$, a contradiction.

It follows that $\pi_1'(s_\theta|\theta) = 1$ for all θ and $\pi_2'(a_s|s,m) = 1$ for all $s \in S^{\text{on}}$ and $m \in M$ such that $\pi_1'(s,m|\theta) > 0$ for some $\theta \in \Theta$. Thus, $\pi^{\delta,\gamma_1,\gamma_2}$ induces the same distribution over $\Theta \times S \times A$ as π in the iterated limit where first $\gamma_1 \to 1$ then $\delta \to 1$ then $\gamma_2 \to 1$. Moreover, since $\pi_1'(s_\theta|\theta) = 1$ for all $\theta \in \Theta$, $\pi_2^{\delta,\gamma_1,\gamma_2} = \phi(\pi_1^{\delta,\gamma_1,\gamma_2}, \mathscr{R}_2^{\gamma_2}(\pi_1^{\delta,\gamma_1,\gamma_2})) = \mathscr{R}_2^{\gamma_2}(\pi_1^{\delta,\gamma_1,\gamma_2})$ in the iterated limit. Thus, $\pi^{\delta,\gamma_1,\gamma_2}$ is a fixed point of $\mathscr{R}^{\delta,\gamma_1,\gamma_2}$ in the iterated limit, which means that π' is a stable profile. ■

OA.5 Proof of Lemma A3

Lemma A3. *If π is a PBE-H that satisfies NWBR, then, for every $s \in S$, either*

1. *$\Theta^\ddagger(s,\pi) \neq \emptyset$, or*
2. *$u_1(\theta,s,a) < u_1(\theta,\pi)$ for all $\theta \in \Theta$ and $a \in BR(\Theta,s)$.*

Proof. Let π be a PBE-H that satisfies NWBR. Fix $s \in S$ and suppose that $\Theta^\ddagger(s,\pi) = \emptyset$. Let $\mathcal{A}_- = \{\alpha \in MBR(\Theta,s) : u_1(\theta,s,\alpha) < u_1(\theta,\pi)\ \forall\theta \in \Theta\}$ be the set of receiver mixed best responses that make playing s strictly worse for every type than their outcome under π. Similarly, let $\mathcal{A}_+ = \{\alpha \in MBR(\Theta,s) : \exists\theta \in \Theta \text{ s.t. } u_1(\theta,s,\alpha) > u_1(\theta,\pi)\}$ be the set of receiver mixed best responses that make some type strictly better off by playing s than receiving their outcome under π. $\mathcal{A}_-$ and $\mathcal{A}_+$ are disjoint open subsets of $MBR(\Theta,s)$, and $\mathcal{A}_- \cup \mathcal{A}_+ = MBR(\Theta,s)$ since $\Theta^\ddagger(s,\pi) = \emptyset$. As $MBR(\Theta,s)$ is connected, either $\mathcal{A}_- = MBR(\Theta,s)$ or $\mathcal{A}_+ = MBR(\Theta,s)$. $\mathcal{A}_+ = MBR(\Theta,s)$ is not possible when π is a PBE-H that satisfies NWBR since then, for every $\alpha \in MBR(\widehat{\Theta}(s,\pi),s)$, there is some θ such that $u_1(\theta,s,\alpha) > u_1(\theta,\pi)$. Therefore, $\mathcal{A}_- = MBR(\Theta,s)$, which gives $u_1(\theta,s,a) < u_1(\theta,\pi)$ for all $a \in BR(\Theta,s)$. ■

OA.6 Omitted Analysis of Examples

OA.6.1 Analysis of Example 2

Proposition OA 4. *The game in Example 2 has stable profiles where all types play Pass with probability* 1.

Proof. We specify that the worker prior g_2 is a Dirichlet distribution. For $m \in \{m_{Hire,\theta_H}, m_{Hire,\{\theta_H,\theta_M\}}\}$, it has initial weight 1 on $(\theta_H, Hire, m)$, 1/2 on $(\theta_M, Hire, m)$, and 1/4 on $(\theta_L, Hire, m)$. For $m = m_{Hire,\theta_M}$, it has initial weight 3/5 on $(\theta_H, Hire, m)$, 1 on $(\theta_M, Hire, m)$, and 1/4 on $(\theta_L, Hire, m)$. For all other messages m, it has initial weight 1/4 on $(\theta_H, Hire, m)$, 1/4 on $(\theta_M, Hire, m)$, and 1 on $(\theta_L, Hire, m)$. Note that initial trust is satisfied: For instance, when a worker first encounters a firm who plays $(Hire, m_{In,\theta_H})$, the probability they place on the firm having type θ_H is 4/7, θ_M is 2/7, and θ_L is 1/7, so $e_H = BR(\theta_H, Hire)$ is optimal.

We observe that e_L is the worker's unique best response to $Hire$ under any distribution that puts probability strictly higher than 3/7 on θ_L. Additionally, if a worker has encountered past play of $(Hire, m)$ and all such plays have been by firms with type θ_L, then the worker will respond to the next instance of $(Hire, m)$ with e_L. To see that this holds for the case $m = m_{Hire,\theta_H}$, note that the worker's conditional distribution over the firm's type after $(Hire, m_{Hire,\theta_H})$ must put probability at least 5/11 on θ_L. Analogous arguments handle the other cases.

We focus on steady state profiles in which, for every $m \in M$, the aggregate probability that a worker responds to $(Hire, m)$ with e_M is less than 1/4. Under such responses, whenever it is weakly optimal for θ_H or θ_M to play $Hire$, it must be strictly optimal for θ_L to do so. To see this, note that

$$u_1(\theta_H, Hire, \alpha) = 21\alpha[e_H] + 6\alpha[e_M] - 5,$$

so $\alpha[e_H] \geq 5/21 - 6/21\alpha[e_M]$ whenever $u_1(\theta_H, Hire, \alpha) \geq 0$, and

$$u_1(\theta_M, Hire, \alpha) = 12\alpha[e_H] + 10\alpha[e_M] - 4,$$

so $\alpha[e_H] \geq 1/3 - 5/6\alpha[e_M]$ whenever $u_1(\theta_M, Hire, \alpha) \geq 0$. Additionally,

$$u_1(\theta_L, Hire, \alpha) = 5\alpha[e_H] + 2\alpha[e_M] - 1,$$

which is strictly positive whenever $\alpha[e_H] \geq \min\{5/21 - 6/21\alpha[e_M], 1/3 - 5/6\alpha[e_M]\}$ and $\alpha[e_M] \leq 1/4$. We argue that such steady state profiles exist in the iterated limit where $\gamma_1 \to 1$, then $\delta \to 1$, and then $\gamma_2 \to 1$, and that the corresponding aggregate probability that any type plays $Hire$ converges to 0.

Let $\chi : \Delta(A) \rightrightarrows \Delta(A)$ be the correspondence given by

$$\chi(\alpha) = \begin{cases} \{\alpha\} & \text{if } \alpha[e_M] < \frac{1}{4} \\ \{\alpha' \in \Delta(A) : \alpha'[e_M] = \frac{1}{4}\} & \text{if } \alpha[e_M] \geq \frac{1}{4} \end{cases},$$

and let $\rho : \Pi_2 \rightrightarrows \Pi_2$ be the correspondence given by

$$\rho(\pi_2) = \{\pi_2' \in \Pi_2 : \pi_2'(\cdot|Hire, m) \in \chi(\pi_2(\cdot|Hire, m)) \ \forall m \in M\}.$$

Note that ρ is upper hemicontinuous, convex-valued, and coincides with the identity correspondence whenever $\pi_2(e_M|In, m) < 1/4$ for all m. Let $\upsilon : \Pi_1 \rightrightarrows \Pi_1$ be the correspondence given by

$$\begin{aligned} \upsilon(\pi_1) = \Bigg\{\pi_1' \in \Pi_1 : &(1)\ \pi_1'[Hire, m|\theta] = \min\left\{\pi_1[Hire, m|\theta], \frac{\lambda(\theta_L)}{2\lambda(\theta)}\right\} \ \forall m \in M,\ \theta \in \{\theta_H, \theta_M\}, \\ &(2)\ \pi_1'[Pass, m|\theta] = \pi_1[Hire, m|\theta] \ \forall m \neq m_{Pass,\theta_H},\ \theta \in \{\theta_H, \theta_M\}, \\ &(3)\ \pi_1'[s, m|\theta_L] = \pi_1[s, m|\theta_L] \ \forall s \in \{Hire, Pass\},\ m \in M, \Bigg\}. \end{aligned}$$

Note that υ is upper hemicontinuous, convex-valued, and coincides with the identity correspondence whenever $\pi_1(Hire, m|\theta) < \lambda(\theta_L)/(2\lambda(\theta))$ for all $m \in M$ and $\theta \in \{\theta_H, \theta_M\}$.

Consider the correspondence $\widetilde{\mathscr{R}}^{\delta,\gamma_1,\gamma_2} : \Pi_1 \times \Pi_2 \rightrightarrows \Pi_1 \times \Pi_2$ given by $\widetilde{\mathscr{R}}^{\delta,\gamma_1,\gamma_2}(\pi_1, \pi_2) = \{(\pi_1', \pi_2') \in \Pi_1 \times \Pi_2 : \pi_1' = \upsilon(\mathscr{R}_1^{\delta,\gamma_1}(\pi_2)) \text{ and } \pi_2' \in \rho(\mathscr{R}_2^{\gamma_2}(\pi_1))\}$. Since $\widetilde{\mathscr{R}}^{\delta,\gamma_1,\gamma_2}$ is upper hemicontinuous and convex-valued, Kakutani's fixed point theorem guarantees the existence of a fixed point $(\pi_1^{\delta,\gamma_1,\gamma_2}, \pi_2^{\delta,\gamma_1,\gamma_2})$.

We establish that $\lim_{\gamma_2\to 1}\lim_{\delta\to 1}\lim_{\gamma_1\to 1}\pi_1^{\delta,\gamma_1,\gamma_2}[Hire|\theta] = 0$ for $\theta \in \{\theta_H, \theta_M\}$. Suppose towards a contradiction that there is a sequence of worker continuation probabilities $\{\gamma_{2,j}\}_{j\in\mathbb{N}}$, a collection of sequences of firm discount factors $\{\delta_{j,k}\}_{j,k\in\mathbb{N}}$, and a collection of sequences of firm continuation probabilities $\{\gamma_{1,j,k,l}\}_{j,k,l\in\mathbb{N}}$ such that (a) $\lim_{j\to\infty}\gamma_{2,j} = 1$, (b) $\lim_{k\to\infty}\delta_{j,k} = 1$ for all j, (c) $\lim_{l\to\infty}\gamma_{1,j,k,l} = 1$ for all j, k, (d) $\lim_{j\to\infty}\lim_{k\to\infty}\lim_{l\to\infty}\pi_1^{\delta_{j,k},\gamma_{1,j,k,l},\gamma_{2,j}}[Hire, m|\theta]$ exists for all $\theta \in \Theta$ and $m \in M$, and (e) $\lim_{j\to\infty}\lim_{k\to\infty}\lim_{l\to\infty}\pi_1^{\delta_{j,k},\gamma_{1,j,k,l},\gamma_{2,j}}[Hire|\theta] > 0$ for either $\theta = \theta_H$ or $\theta = \theta_M$. Then since $\pi_2^{\delta,\gamma_1,\gamma_2}(e_M|Hire, m) \leq 1/4$ for all $m \in M$, Lemma B1 implies that $\lim_{j\to\infty}\lim_{k\to\infty}\lim_{l\to\infty}\pi_1^{\delta_{j,k},\gamma_{1,j,k,l},\gamma_{2,j}}[Hire|\theta_L] = 1$. Therefore, there exists some $m \in M$ such that $\lim_{j\to\infty}\lim_{k\to\infty}\lim_{l\to\infty}\pi_1^{\gamma_{1,j,k,l},\gamma_{2,j}}[Hire, m|\theta_L] > 0$ and

$$\lim_{j\to\infty}\lim_{k\to\infty}\lim_{l\to\infty}\pi_1^{\delta_{j,k},\gamma_{1,j,k,l},\gamma_{2,j}}[Hire|\theta_L] \geq \frac{\lambda(\theta_L)}{2\lambda(\theta)}\lim_{j\to\infty}\lim_{k\to\infty}\lim_{l\to\infty}\pi_1^{\delta_{j,k},\gamma_{1,j,k,l},\gamma_{2,j}}[Hire|\theta]$$

for both $\theta \in \{\theta_H, \theta_M\}$. By Lemma 2 and the fact that the unique worker best response to $Hire$ is e_L when the probability the type is θ_L is at least 1/2, this implies that $\lim_{j\to\infty}\lim_{k\to\infty}\lim_{l\to\infty}\mathscr{R}_2^{\gamma_{2,j}}(\pi_1^{\delta_{j,k},\gamma_{1,j,k,l},\gamma_{2,j}})(e_L|Hire, m) = 1$. Since $\chi(\pi_2(\cdot|Hire, m)) = \{\pi_2(\cdot|Hire, m)\}$ if $\pi_2(e_M|Hire, m) < 1/4$, it follows that $\lim_{j\to\infty}\lim_{k\to\infty}\lim_{l\to\infty}\pi_2^{\delta_{j,k},\gamma_{1,j,k,l},\gamma_{2,j}}(\pi_1)(e_L|Hire, m) = 1$. However, by Lemma B1, this requires that $\lim_{j\to\infty}\lim_{k\to\infty}\lim_{l\to\infty}\pi_1^{\delta_{j,k},\gamma_{1,j,k,l},\gamma_{2,j}}[Hire, m] = 0$ must hold, a contradiction.

A similar argument establishes that $\lim_{\gamma_2\to 1}\lim_{\delta\to 1}\lim_{\gamma_1\to 1}\pi_1^{\delta,\gamma_1,\gamma_2}[Hire|\theta_L] = 0$, so $\lim_{\gamma_2\to 1}\lim_{\delta\to 1}\lim_{\gamma_1\to 1}\pi_1^{\delta,\gamma_1,\gamma_2}[Hire] = 0$. Since a worker will only play e_M in response to some $(Hire, m)$ if they have previously encountered a firm playing $(Hire, m)$, we have that $\mathscr{R}_2^{\gamma_{2,j}}(\pi_1^{\gamma_{1,k,l},\gamma_{2,k}})(e_M|Hire,m) < 1/4$ for all $m \in M$ in the iterated limit. Since $\rho(\pi_2) = \{\pi_2\}$ if $\pi_2(e_M|Hire,m) < 1/4$ for all m, $\pi_2^{\delta,\gamma_1,\gamma_2} = \rho(\mathscr{R}_2^{\gamma_2}(\pi_1^{\delta,\gamma_1,\gamma_2})) = \mathscr{R}_2^{\gamma_2}(\pi_1^{\delta,\gamma_1,\gamma_2})$ for fixed, sufficiently high $\gamma_2 \in [0,1)$ when δ is sufficiently close to 1 and, given δ, γ_1 is sufficiently close to 1. For similar reasons, $\pi_1^{\delta,\gamma_1,\gamma_2} = \upsilon(\mathscr{R}_1^{\delta,\gamma_1}(\pi_2^{\delta,\gamma_1,\gamma_2})) = \mathscr{R}_2^{\gamma_2}(\pi_1^{\delta,\gamma_1,\gamma_2})$ also holds in the iterated limit. Thus, $(\pi_1^{\delta,\gamma_1,\gamma_2}, \pi_2^{\delta,\gamma_1,\gamma_2})$ is a fixed point of $\mathscr{R}^{\delta,\gamma_1,\gamma_2}$ for fixed, sufficiently high $\gamma_2 \in [0,1)$, when δ is sufficiently close to 1 and, given δ, γ_1 is sufficiently close to 1. We conclude that there are stable profiles in which every type plays *Pass*. ■

OA.6.2 Analysis of Example 3

Proposition OA 5. *The game in Example 3 has stable profiles where both types play Out with probability* 1*.*

Proof. We specify that the receiver prior g_2 is a Dirichlet distribution with initial weight 1 on $(\theta_1, In, m_{In,\theta_1})$ and 1/2 on $(\theta_2, In, m_{In,\theta_1})$, and, for all other messages $m \neq m_{In,\theta_1}$, initial weight 1/2 on (θ_1, In, m) and 1 on (θ_2, In, m). This means that initial trust is satisfied: When a receiver first encounters a sender who plays $(In, m_{In,\theta})$, the probability they place on the receiver having type θ is 2/3 so $BR(\theta, In)$ is optimal.

We claim first that if a receiver has encountered past plays of (In, m) and all such plays have been by senders with the same type θ, then the receiver will respond to the next instance of (In, m) with $BR(\theta, In)$. We demonstrate this for the case $m = m_{\theta_1}$; analogous arguments handle the other case. If this message has only ever been sent by θ_1, the receiver's belief about the sender's type after (In, m_{θ_1}) must put probability at least $(1+1)/(1+1+.5) = 4/5$ on θ_1, which makes a_1 the unique receiver best response. When $\theta = \theta_2$, the receiver's conditional distribution over the sender's type after (In, m_{θ_1}) must put probability at least $(1+.5)/(1+1+.5) = 3/5$ on θ_2, which makes a_2 the unique receiver best response.

We focus on steady state profiles in which, for every $m \in M$, the aggregate probability that a receiver responds to (In, m) with a_3 is less than 1/4. Under such responses, it can never be weakly optimal for both types to play In with the same message. To see this, note that

$$u_1(\theta_1, In, \alpha) + u_2(\theta_2, In, \alpha) = -\alpha[a_1] - \alpha[a_2] + 2\alpha[a_3] = -1 + 3\alpha[a_3],$$

which is strictly negative whenever $\alpha[a_3] \leq 1/4$. We argue that such steady state profiles exist in the iterated limit where $\gamma_1 \to 1$ then $\delta \to 1$ then $\gamma_2 \to 1$ and that the corresponding aggregate probability that either sender type plays In converges to 0.

Let $\chi : \Delta(A) \rightrightarrows \Delta(A)$ be the correspondence given by

$$\chi(\alpha) = \begin{cases} \{\alpha\} & \text{if } \alpha[a_3] < \frac{1}{4} \\ \{\alpha' \in \Delta(A) : \alpha'[a_3] = \frac{1}{4}\} & \text{if } \alpha[a_3] \geq \frac{1}{4} \end{cases},$$

and let $\rho : \Pi_2 \rightrightarrows \Pi_2$ be the correspondence given by

$$\rho(\pi_2) = \{\pi_2' \in \Pi_2 : \pi_2'(\cdot|In, m) \in \chi(\pi_2(\cdot|In, m)) \ \forall m \in M\}.$$

Note that ρ is upper hemicontinuous, convex-valued, and coincides with the identity correspondence whenever $\pi_2(a_3|In, m) < 1/4$ for all m.

Consider the correspondence $\widetilde{\mathscr{R}}^{\delta,\gamma_1,\gamma_2} : \Pi_1 \times \Pi_2 \rightrightarrows \Pi_1 \times \Pi_2$ given by $\widetilde{\mathscr{R}}^{\delta,\gamma_1,\gamma_2}(\pi_1, \pi_2) = \{(\pi_1', \pi_2') \in \Pi_1 \times \Pi_2 : \pi_1' = \mathscr{R}_1^{\delta,\gamma_1}(\pi_2) \text{ and } \pi_2' \in \rho(\mathscr{R}_2^{\gamma_2}(\pi_1))\}$. Since $\mathscr{R}$ is upper hemicontinuous and convex-valued, Kakutani's fixed point theorem guarantees the existence of a fixed point $(\pi_1^{\delta,\gamma_1,\gamma_2}, \pi_2^{\delta,\gamma_1,\gamma_2})$. As $\pi_2^{\delta,\gamma_1,\gamma_2}(a_3|s, m) \leq 1/4$ for all (s, m) by construction, Lemma B1 implies that, for all $\gamma_2 \in [0, 1)$ and (s, m), either $\lim_{\gamma_1 \to 1} \pi_1^{\delta,\gamma_1,\gamma_2}[In, m|\theta_1] = 0$ or $\lim_{\gamma_1 \to 1} \pi_1^{\delta,\gamma_1,\gamma_2}[In, m|\theta_2] = 0$. This means that, as $\gamma_1 \to 1$ then $\delta \to 1$, the probability that a receiver encounters senders with both types that pair In with the same message m approaches 0. Since a receiver would only ever play a_3 in response to (In, m) if they have previously encountered senders of both types play (In, m), this means that $\lim_{\delta \to 1} \lim_{\gamma_1 \to 1} \mathscr{R}_2^{\gamma_2}(\pi_1^{\delta,\gamma_1,\gamma_2})(a_3|In, m) = 0$ for all $m \in M$. Since $\rho(\pi_2) = \{\pi_2\}$ if $\pi_2(a_3|In, m) < 1/4$ for all m, $\pi_2^{\delta,\gamma_1,\gamma_2} = \rho(\mathscr{R}_2^{\gamma_2})(\pi_1^{\delta,\gamma_1,\gamma_2}) = \mathscr{R}_2^{\gamma_2}(\pi_1^{\delta,\gamma_1,\gamma_2})$ for fixed

$\gamma_2 \in [0,1)$ when δ is sufficiently close to 1 and, given δ, γ_1 is sufficiently close to 1. Thus, for fixed $\gamma_2 \in [0,1)$, $(\pi_1^{\delta,\gamma_1,\gamma_2}, \pi_2^{\delta,\gamma_1,\gamma_2})$ is a fixed point of $\mathscr{R}^{\delta,\gamma_1,\gamma_2}$ when δ is sufficiently close to 1 and, given δ, γ_1 sufficiently close to 1.

To show that $\lim_{\gamma_2\to 1}\lim_{\delta\to 1}\lim_{\gamma_1\to 1}\pi_1^{\delta,\gamma_1,\gamma_2}[In] = 0$, suppose towards a contradiction that there is a sequence of receiver continuation probabilities $\{\gamma_{2,j}\}_{j\in\mathbb{N}}$, a collection of sequences of sender discount factors $\{\delta_{j,k}\}_{j,k\in\mathbb{N}}$, and a collection of sequences of sender continuation probabilities $\{\gamma_{1,j,k,l}\}_{j,k,l\in\mathbb{N}}$ such that (a) $\lim_{j\to\infty}\gamma_{2,j} = 1$, (b) $\lim_{k\to\infty}\delta_{j,k} = 1$ for all j, (c) $\lim_{l\to\infty}\gamma_{1,j,k,l} = 1$ for all j,k, and (d) $\lim_{j\to\infty}\lim_{k\to\infty}\lim_{l\to\infty}\pi_1^{\delta_{j,k},\gamma_{1,j,k,l},\gamma_{2,j}}[In, m|\theta] > 0$ for some $\theta \in \Theta$ and $m \in M$. Without loss of generality, take $\theta = \theta_1$. By what we have shown, it must be that $\lim_{k\to\infty}\lim_{l\to\infty}\pi_1^{\delta_{j,k},\gamma_{1,j,k,l},\gamma_{2,j}}[In, m|\theta_2] = 0$ for all sufficiently large j. Combining this with $\lim_{j\to\infty}\lim_{k\to\infty}\lim_{l\to\infty}\pi_1^{\delta_{j,k},\gamma_{1,j,k,l},\gamma_{2,j}}[In, m|\theta_1] > 0$ and $\lim_{j\to\infty}\gamma_{2,j} = 1$ gives $\lim_{j\to\infty}\lim_{k\to\infty}\lim_{l\to\infty}\pi_2^{\delta_{j,k},\gamma_{1,j,k,l},\gamma_{2,j}}(a_1|s,m) = 1$, because with probability 1 every receiver encounters a type θ_1 sender playing (In, m) but never encounters a type θ_2 sender playing (In, m). However, since $u_1(\theta_1, In, a_1) < 0$, $\lim_{j\to\infty}\lim_{k\to\infty}\lim_{l\to\infty}\pi_2^{\delta_{j,k},\gamma_{1,j,k,l},\gamma_{2,j}}(a_1|s,m) = 1$ combined with Lemma B1 requires $\lim_{j\to\infty}\lim_{k\to\infty}\lim_{l\to\infty}\pi_1^{\delta_{j,k},\gamma_{1,j,k,l},\gamma_{2,j}}[In, m|\theta_1] = 0$, a contradiction. ■

OA.6.3 Analysis of Example 4

Proposition OA 6. *The least-cost separating equilibrium of the game in Example 4 has $\theta = 1$ play $(s_1^*(1), s_2^*(1)) = (1/2, 0)$, to which the receiver responds $a^*(1) = 10$, $\theta = 2$ play $(s_1^*(2), s_2^*(2)) = (1/2, 5)$, to which the receiver responds $a^*(2) = 20$, and $\theta = 3$ play $(s_1^*(3), s_2^*(3)) = (1/2, 15)$, to which the receiver responds $a^*(3) = 30$.*

Proof. We first establish that this play is consistent with a separating PBE. Given an arbitrary (s_1, s_2) and a belief $\tilde{\lambda}$ about the sender's type, the receiver's best responses are the closest actions to $20s_1\mathbb{E}_{\tilde{\lambda}}[\theta]$, as can be readily verified using the receiver's utility function. For $s_1 = 1/2$ and the belief that the type is θ, the receiver's best response is 10θ, so the prescribed receiver play following the on-path sender play is indeed optimal.

Fix the receiver's response to any off-path signal-message pair (s_1, s_2, m) to be $20s_1$, i.e. the best response under a belief putting probability 1 on $\theta = 1$. All that remains is to check that the incentives of the sender types are satisfied. We verify this for the $\theta = 3$ sender type. (Similar arguments handle the other two types.) Under the prescribed play, the payoff of the $\theta = 3$ sender type is $u_1(3, 1/2, 15, 30) = 30$. If the $\theta = 3$ sender were instead to mimic $\theta = 1$ or $\theta = 3$, their payoff would be 15 or 25, respectively. Moreover, if the $\theta = 3$ sender were to deviate to some off-path signal-message pair (s_1, s_2, m), their payoff would be $60(1 - s_1)s_1 - s_2$, which is strictly lower than 30 for all $s_1 \in [0, 1]$ and $s_2 \geq 0$.

We now show that every other separating equilibrium results in (weakly) lower payoffs to each of the sender types. The payoff of the $\theta = 1$ sender from (s_1, s_2) when the receiver responds with $20s_1$ is $20(1 - s_1)s_1 - s_2$, which attains its maximum value of 5 at $(s_1^*(1), s_2^*(1))$. The maximum possible payoff of the $\theta = 2$ sender from playing some (s_1, s_2) when the receiver responds with $40s_1$, subject to the constraint that $\theta = 1$ would obtain a lower payoff than 5 by imitating $\theta = 2$ is

$$\max_{(s_1,s_2)\in S} 80(1 - s_1)s_1 - s_2 \text{ s.t. } 40(1 - s_1)s_1 - s_2 \leq 5.$$

The solution to this problem is $(s_1^*(2), s_2^*(2))$, and the resulting payoff to $\theta = 2$ is 15. Finally, the maximum possible payoff of the $\theta = 3$ sender from playing some (s_1, s_2) when the receiver responds with $60s_1$, subject to the constraint that $\theta = 2$ would obtain a lower payoff than 15 by imitating $\theta = 3$ is

$$\max_{(s_1,s_2)\in S} 120(1 - s_1)s_1 - s_2 \text{ s.t. } 80(1 - s_1)s_1 - s_2 \leq 15.$$

The solution to this problem is $(s_1^*(3), s_2^*(3))$. ■

Proposition OA 7. *If π is a JCE in the game in Example 4, then each θ plays $(s_1^*(\theta), s_2^*(\theta))$ with strictly positive probability, and the receiver responds to all on-path $(s_1^*(\theta), s_2^*(\theta), m)$ with $a^*(\theta)$ as in the least-cost separating equilibrium.*

Proof. We first establish that in a JCE π, for each signal-message pair (s_1, s_2, m) played by $\theta = 3$, the product of $(1 - s_1)$ and the receiver's response has expected value at least $44/3$. Suppose otherwise that there is some signal-message pair (s_1, s_2, m) that $\theta = 3$ plays which induces a receiver response with expected value $\tilde{a}$ such that $(1 - s_1)\tilde{a} < 44/3$. It must be that $s_2 < 44$, as otherwise $\theta = 3$ would obtain a strictly negative payoff. Thus, $s_2' = \lceil s_2 + 30 - 2(1 - s_1)\tilde{a} \rceil \in S$. Note that $u_1(3, \pi) = 3(1 - s_1)\tilde{a} - s_2$, while $u_1(\theta, \pi) \leq \theta(1 - s_1)\tilde{a} - s_2$ for $\theta \in \{1, 2\}$. Since $u_1(3, 1/2, s_2', a) = 3a/2 - s_2'$, we have that $u_1(3, 1/2, s_2', a) \geq u_1(3, \pi)$ if and only if $a \geq 2(1 - s_1)\tilde{a} + 2(s_2' - s_2)/3$, with the inequality strict for all $a > 2(1 - s_1)\tilde{a} + 2(s_2' - s_2)/3$. Moreover, $u_1(\theta, 1/2, s_2', a) \geq u_1(\theta, \pi)$ for $\theta = 1$ or $\theta = 2$ only if $u_1(\theta, 1/2, s_2', a) = \theta a/2 - s_2' \geq \theta(1 - s_1)\tilde{a} - s_2$, which requires $a \geq 2(1 - s_1)\tilde{a} + s_2' - s_2$. Since $s_2' > s_2$, $2(1 - s_1)\tilde{a} + s_2' - s_2 > 2(1 - s_1)\tilde{a} + 2(s_2' - s_2)/3$ which means that $\overline{\Theta}(1/2, s_2', \pi) = \{3\}$ and the only justified response to $(1/2, s_2')$ is 30. As this is strictly greater than $2(1 - s_1)\tilde{a} + 2s_2'/3 - 2s_2/3$ when $(1 - s_1)\tilde{a} < 44/3$, the claim follows.

An immediate implication is that there must be some signal-message pair that $\theta = 2$ sends with positive probability that $\theta = 3$ does not send, because $(1 - s_1)a \leq 25/2$ for any signal (s_1, s_2) and receiver best response a to a belief where the relative weight on $\theta = 2$ versus $\theta = 3$ is at least that of the prior.

We now show that, for each signal-message pair (s_1, s_2, m) played by $\theta = 2$ but not by $\theta = 3$, the product of $1 - s_1$ and the receiver's response must have an expected value between $19/2$ and 10. Whenever the probability of $\theta = 3$ is 0, the product of $(1 - s_1)$ and any undominated receiver response is no more than 10, so we need only show that the expected value of the product must exceed $19/2$. Suppose otherwise that there is some signal-message pair (s_1, s_2, m) that $\theta = 2$ plays but $\theta = 3$ does not play for which the expected value of the receiver response $\tilde{a}$ satisfies $(1 - s_1)\tilde{a} < 19/2$. It must be that $s_2 < 19$, so $s_2' = \lceil s_2 + 10 - (1 - s_1)a \rceil \in S$. Note that $u_1(2, \pi) = 2(1 - s_1)\tilde{a} - s_2$, while $u_1(1, \pi) \leq (1 - s_1)\tilde{a} - s$. Since $u_1(2, 1/2, s_2', a) = a - s_2'$, we have that $u_1(2, 1/2, s_2', a) \geq u_1(2, \pi)$ if and only if $a \geq 2(1 - s_1)\tilde{a} + s_2' - s_2$, with the inequality strict for all $a > 2(1 - s_1)\tilde{a} + s_2' - s_2$. Moreover, $u_1(1, 1/2, s_2', a) \geq u_1(1, \pi)$

only if $a/2 - s_2' \geq (1-s_1)\tilde{a} - s$, which requires $a \geq 2(1-s_1)\tilde{a} + 2(s_2' - s_2)$. Since $s_2' > s_2$, $2(1-s_1)\tilde{a} + 2(s_2' - s_2) > 2(1-s_1)\tilde{a} + s_2' - s_2$, which means that $\overline{\Theta}(s+10, \pi) \subseteq \{2, 3\}$ so justified responses to $(1/2, s_2')$ must weakly exceed 20. As this is strictly greater than $2(1-s_1)\tilde{a} + s_2' - s_2$ when $(1-s_1)\tilde{a} < 19/2$, the claim follows.

There must be some signal-message pair that only $\theta = 1$ plays. To see this, first observe that there can be no signal-message pair played by both $\theta = 1$ and $\theta = 3$. If there were some signal-message pair (s_1, s_2, m) played by both $\theta = 1$ and $\theta = 3$, the product of $1 - s_1$ and the expected value of the receiver response $\tilde{a}$ must be less than $25/2$, because increasing differences in θ and $(1-s_1)a$ in the sender utility function implies that every signal-message pair played by $\theta = 2$ must induce the same expected value $(1-s_1)\tilde{a}$. This contradicts the fact that, for every signal-message pair played by $\theta = 3$, the product of $1 - s_1$ and the expected value of the receiver response must be weakly greater than $44/3$. Additionally, $\theta = 1$ cannot only play signal-message pairs that are also played by $\theta = 2$. Otherwise, there would be some signal-message pair (s_1, s_2, m) played by $\theta = 2$, for which the product of $1 - s_1$ and the receiver response would have expected value weakly less than $15/2$ since $(1-s_1)a \leq 15/2$ for any receiver best response a to a belief where the weight on $\theta = 3$ is 0 and the weight on $\theta = 1$ is at least that of the prior.

For every signal-message pair that only $\theta = 1$ plays, $s_1 = 1/2$, $s_2 = 0$, and the receiver responds with $a = 10$. The reason is the payoff $\theta = 1$ obtains from a signal-message pair (s_1, s_2, m) that only $\theta = 1$ plays is $20(1-s_1)s_1 - s_2$, which is strictly less than 5 if $s_1 \neq 1/2$ or $s_2 > 0$. However, $\theta = 1$ can secure a payoff of 5 by simply playing $(s_1, s_2) = (1/2, 0)$, since every $a < 10$ is a strictly dominated response for the receiver.

We now argue that, for every signal-message pair played by $\theta = 2$ but not by $\theta = 3$, $s_1 = 1/2$, $s_2 = 5$, and the receiver responds with $a = 20$. We have previously established that the product of $1 - s_1$ and the expected value of the receiver's response $\tilde{a}$ must be between $19/2$ and 10. For $(1-s_1)\tilde{a} < 10$ to hold, it must be that $\theta = 1$ also plays this signal-message pair. This requires $u_1(1, s, \tilde{a}) = (1-s_1)\tilde{a} - s_2 = u_1(1, \pi)$. As previously established, $u_1(1, \pi) = 5$, so it must be that $s_2 = (1-s_1)\tilde{a} - 5$. However, there is no $\tilde{a} \in$

$[19/2, 10)$ such that $\tilde{a} - 5 \in S$. Therefore, $(1 - s_1)\tilde{a} = 10$. Since $(1 - s_1)\tilde{a} \leq 40(1 - s_1)s_1$ and $40(1 - s_1)s_1 < 10$ for all $s_1 \neq 1/2$, it follows that $s_1 = 1/2$ and thus $\tilde{a} = 20$. From $u_1(1, 1/2, s_2, 20) = 10 - s_2 \leq 5 = u_1(1, \pi)$, we obtain $s_2 \geq 5$. All that remains is to rule out $s_2 > 5$. If $s_2 > 5$, $u_1(1, 1/2, s_2 - 1, a) = a/2 - s_2 + 1 \geq 5 = u_1(1, \pi)$ only if $a \geq 20$. On the other hand, $u_1(2, 1/2, s_2 - 1, a) = a - s_2 + 1 \geq 20 - s_2 = u_1(2, \pi)$ if and only if $a \geq 19$, with the inequality strict for all $a > 19$. Thus, $\overline{\Theta}(1/2, s_2 - 1, \pi) \subseteq \{2, 3\}$, so justified responses to $(1/2, s_2 - 1)$ must weakly exceed 20. It follows that $s_2 = 5$.

Finally, we show that, for every signal-message pair played by $\theta = 3$, $s_1 = 1/2$, $s_2 = 15$, and the receiver responds with $a = 40$. We have previously established that the product of $1 - s_1$ and the expected value of the receiver's response $\tilde{a}$ must be between 44/3 and 15. For $(1 - s_1)\tilde{a} < 15$ to hold, it must be that $\theta = 2$ also plays this signal-message pair. This requires $u_1(2, s_1, s_2, \tilde{a}) = 2(1 - s_1)\tilde{a} - s_2 = u_1(2, \pi)$. As previously established, $u_1(2, \pi) = 15$, so it must be that $s_2 = 2(1 - s_1)\tilde{a} - 15$. However, there is no $(1 - s_1)\tilde{a} \in [44/3, 15)$ such that $2(1 - s_1)\tilde{a} - 15 \in S$. Therefore, $(1 - s_1)\tilde{a} = 15$. Since $(1 - s_1)\tilde{a} \leq 60(1 - s_1)s_1$ and $60(1 - s_1)s_1 < 15$ for all $s_1 \neq 1/2$, it follows that $s_1 = 1/2$ and thus $\tilde{a} = 30$. From $u_1(2, 1/2, s_2, 30) = 30 - s_2 \leq 15 = u_1(2, \pi)$, we obtain $s_2 \geq 15$. All that remains is to rule out $s > 15$. If $s > 15$, $u_1(\theta, 1/2, s_2 - 1, a) = \theta a/2 - s + 1 \geq u_1(\theta, \pi)$ for either $\theta = 1$ or $\theta = 2$ requires that $a \geq 40$. On the other hand, $u_1(3, 1/2, s_2 - 1, a) = 3a/2 - s_2 + 1 \geq 45 - s_2 = u_1(3, \pi)$ if and only if $a \geq 29/3$, with the inequality strict for all $a > 29/3$. Thus, $\overline{\Theta}(1/2, s_2 - 1, \pi) = \{3\}$, so the only justified response to $(1/2, s_2 - 1)$ is 30. It follows that $s_2 = 15$. ■

OA.7 Other Examples

OA.7.1 Stability without Initially Trusting Receivers

Example OA 1. The sender's type space is $\Theta = \{\theta_1, \theta_2\}$, signal space is $S = \{In, Out\}$, and the receiver's action space is $A = \{a_1, a_2\}$. The payoffs to the sender and receiver are given below.

θ_1	a_1	a_2
In	$1, 1$	$-1, -1$
Out	$0, 0$	$0, 0$

θ_2	a_1	a_2
In	$-1, -1$	$-1, 1$
Out	$0, 0$	$0, 0$

Out strictly dominates In for type θ_2, so θ_2 plays Out in every equilibrium of this game. However, there are equilibria in which θ_1 plays In and equilibria in which θ_1 plays Out. The equilibria where θ_1 plays Out do not survive the Intuitive Criterion since a_1 is the receiver's unique best response to In when the sender's type is θ_1, and θ_1 obtains a strictly higher payoff from (In_1, a_1) than from playing Out.

We show that, when g_2 is such that a receiver plays a_2 when they first encounter a sender playing (In, m) for every message $m \in M$, there are stable profiles in which θ_1 plays Out.

We focus on steady state profiles in which the aggregate probability that a receiver responds to (In, m) with a_1 is less than 1/3 for every message $m \in M$, which makes it strictly optimal for type θ_1 senders to play Out. We show that, for fixed $\gamma_2 \in [0, 1)$, such steady state profiles exist, and, moreover, that the corresponding aggregate probability that a type θ_1 sender plays In approaches 0 as $\gamma_1 \to 1$ and then $\delta \to 1$.

Let $\psi : \Pi_2 \to \Pi_2$ be the mapping given by

$$\psi(\pi_2)(a_1|In, m) = \min\left\{\pi_2(a_1|In, m), \frac{1}{3}\right\} \ \forall m \in M.$$

Note that ψ is continuous and coincides with the identity mapping whenever $\pi_2(a_1|In, m) \leq 1/3$ for all m.

Consider the mapping $\widetilde{\mathscr{R}}^{\delta,\gamma_1,\gamma_2} : \Pi_1 \times \Pi_2 \to \Pi_1 \times \Pi_2$ given by $\widetilde{\mathscr{R}}^{\delta,\gamma_1,\gamma_2}(\pi_1, \pi_2) = (\mathscr{R}_1^{\delta,\gamma_1}(\pi_2), \psi(\mathscr{R}_2^{\gamma_2}(\pi_1)))$. Since $\widetilde{\mathscr{R}}^{\delta,\gamma_1,\gamma_2}$ is continuous, Brouwer's fixed point theorem guarantees the existence of a fixed point $(\pi_1^{\delta,\gamma_1,\gamma_2}, \pi_2^{\delta,\gamma_1,\gamma_2})$. As $\pi_2^{\delta,\gamma_1,\gamma_2}(a_1|In, m) \leq 1/3$ for all m by construction, Lemma B1 implies that $\lim_{\delta\to 1} \lim_{\gamma_1 \to 1} \pi_1^{\delta,\gamma_1,\gamma_2}[In] = 0$ for all $\gamma_2 \in [0, 1)$. Furthermore, because g_2 is such that every receiver would play a_2 at a first encounter with a sender playing (In, m), $\lim_{\delta\to 1}\lim_{\gamma_1\to 1}\mathscr{R}_2^{\gamma_2}(\pi_1^{\delta,\gamma_1,\gamma_2})(a_1|In,m) = 0$ for all m, $\gamma_2 \in [0,1)$, so the $\pi_2(a_1|In,m) \le 1/3$ constraint does not bind when δ is sufficiently close to 1 and, given δ, γ_1 is sufficiently close to 1. Formally, since $\pi_2^{\delta,\gamma_1,\gamma_2} \neq \mathscr{R}_2^{\gamma_2}(\pi_1^{\delta,\gamma_1,\gamma_2})$ only if $\mathscr{R}_2^{\gamma_2}(\pi_1^{\delta,\gamma_1,\gamma_2})(a_1|In,m) > 1/3$ for some m, we have that, for fixed $\gamma_2 \in [0,1)$, $\pi_2^{\delta,\gamma_1,\gamma_2} = \mathscr{R}_2^{\gamma_2}(\pi_1^{\delta,\gamma_1,\gamma_2})$ for δ sufficiently close to 1 and, given δ, γ_1 sufficiently close to 1. Combining this with the fact that $\pi_1^{\delta,\gamma_1,\gamma_2} = \mathscr{R}_1^{\delta,\gamma_1}(\pi_2^{\delta,\gamma_1,\gamma_2})$ for all $\gamma_1, \gamma_2 \in [0,1)$, it follows that, for fixed $\gamma_2 \in [0,1)$, $(\pi_1^{\delta,\gamma_1,\gamma_2}, \pi_2^{\delta,\gamma_1,\gamma_2})$ is a fixed point of $\mathscr{R}^{\delta,\gamma_1,\gamma_2}$ for δ sufficiently close to 1 and, given δ, γ_1 sufficiently close to 1. Since $\lim_{\gamma_2\to 1}\lim_{\delta\to 1}\lim_{\gamma_1\to 1}\pi_1^{\delta,\gamma_1,\gamma_2}[In] = 0$, we conclude that there are stable profiles in which both types plays *Out*. $\square$

In this example, In is strictly dominated for type θ_2. If the priors of the receiver agents put 0 probability on sender types for whom a given signal is strictly dominated after an observation of that signal, the receivers would respond to In with a_1, which would preclude the "All *Out*" equilibria. Depending on the context, such belief restrictions might be plausible, though they do rely on the receivers knowing the sender payoff function. However, even with such restrictions, stability can still allow implausible outcomes when initial trust is not satisfied. For example, we could modify the payoffs above so that In is no longer strictly dominated for θ_2, but rather conditionally dominated when the receiver response to *Out* uses a particular action, say a_2, with high probability. When the receiver priors are non-degenerate, we could choose the receiver payoffs so that both types playing *Out* is stable.[1]

OA.7.2 Alternate Example Where D1 Does Not Imply JCE

Example OA 2. Here we analyze a simple example that is related to the idea of corporate culture as a way of telling workers what to do in unforeseen contingencies (see e.g. Camerer and Vepsalainen (1988) and Kreps (1990)). The sender is a firm, and the

[1] We could further restrict the receiver priors to assign probability 0 to sender types for whom a given signal is equilibrium dominated, but such restrictions are not consistent with a learning foundation for equilibrium, since they require that the receivers know the equilibrium being played in the population.

receiver is a recently hired worker. The firm's signal $s \in \{Creative, Standard\}$ is its choice of job assignment for the worker: The firm can either assign the worker to one of its "standard" jobs or to a "creative" job. Standard jobs carry out the firm's operation as currently designed, and let the firm effectively control the actions of workers through a combination of monitoring and provision of incentives. Creative jobs are intended to lead to innovations which the firm can then incorporate into its main operations, and the firm has relatively little direct control over the work these workers carry out. The worker's choice of action $a \in \{a_1, a_2, a_3\}$ represents the focus and intensity of their costly effort when assigned a creative job: a_1 and a_2 both represent intense effort directed at productive innovation but with focuses in different sectors, while a_3 represents a lack of productive effort.

The firm has three possible types, $\Theta = \{\theta_1, \theta_2, \theta_3\}$. Type θ_1 and θ_2 firms obtain higher payoffs than the relatively unproductive type θ_3 firms. Moreover, type θ_1 firms are particularly well suited to exploit innovations that workers with creative jobs choosing action a_1 may create, and type θ_2 firms have an advantage with innovations from a_2. Due to their high payoffs from standard jobs, type θ_1 gains relatively less from a worker with a creative job working on a_2 than type θ_3 does. (Likewise for type θ_2 and a_1.) A worker with a creative job is incentivized by rewards that come from successful innovation, so such a worker would like to take action a_1 if the firm has type θ_1, a_2 if the firm has type a_2, and a_3 if the firm has type θ_3.

The payoffs are given below.[2]

[2]The table indicates the worker can take any action in $\{a_1, a_2, a_3\}$ when assigned a standard job. However, we think of the firm as controlling the actual effort of a worker with a standard job, which is why the payoffs are independent of the formal action of a worker assigned a standard job.

θ_1	a_1	a_2	a_3
Creative	$4, 1$	$2, 0$	$0, -1$
Standard	$2, 0$	$2, 0$	$2, 0$

θ_2	a_1	a_2	a_3
Creative	$2, 0$	$4, 1$	$0, -1$
Standard	$2, 0$	$2, 0$	$2, 0$

θ_3	a_1	a_2	a_3
Creative	$1, 0$	$1, 0$	$-1, 1$
Standard	$0, 0$	$0, 0$	$0, 0$

In every JCE, there is a positive probability of the worker being assigned a creative job. The reason is that the worker must, with positive probability, respond to *Creative* with a_3 in order to deter the firm from playing *Creative*, but there is no justified response to *Creative* that uses a_3, because a_3 is an optimal response to *Creative* only when the worker assigns a positive probability to the firm being type θ_3. However, either θ_1 or θ_2 strictly prefers to play *Creative* whenever θ_3 weakly prefers *Creative*, so θ_3 is not a justified type for *Creative*.

Every stable profile has a positive probability of the worker being assigned a creative job because, for every firm type to learn that *Standard* is weakly optimal, the aggregate worker response must use a_3 with positive probability whenever *Creative* is played. Since responding to *Creative* with a_3 is optimal only for beliefs with positive probability on θ_3, Initial Trust implies that some θ_3 firms must be learning to play *Creative* while claiming to be either type θ_1 or θ_2. But if θ_3 firms learn that it is weakly optimal to play *Creative*, then either the θ_1 or θ_2 firms learn that it is strictly optimal to do so.

Unlike JCE, many existing refinements allow equilibria in which all types play *Standard*. We discuss why this is the case for D1, which is typically thought of as a strong refinement. D1 allows the worker to respond to *Creative* with a_3, because there is no type which strictly prefers to play *Creative* whenever θ_3 weakly prefers to do so. In particular, θ_3 strictly prefers to play *Creative* whenever the worker plays either a_1 or a_2 with probability 1. For the other two types, there are some mixtures over a_1

and a_2 at which *Creative* is strictly preferred to *Standard* and others where *Standard* is strictly preferred to *Creative*. In contrast, θ_3 is not a justified type for *Creative* because whenever θ_3 weakly prefers to play *Creative*, there is some type that strictly prefers to do so, but this type need not be the same across worker responses. □

OA.8 Stability Under Alternative Assumptions

OA.8.1 Weakening Initial Trust

Here we discuss a refinement satisfied by all stable profiles under an alternative assumption to initial trust. Suppose that receivers know the payoff functions of the senders, as in Fudenberg and He (2020). Then receivers who are long-lived may feel that they have acquired a good sense of each sender type's equilibrium payoff. Suppose that such a receiver encounters a sender playing a pair $(s, m_{s,\widetilde{\Theta}})$ that the receiver has not previously seen types outside of $\widetilde{\Theta}$ play. If the receiver believes that only types in $\widetilde{\Theta}$ could improve their outcome by deviating to s when the receiver's response is contained in $BR(s, \widetilde{\Theta})$, we assume the receiver finds such a message credible and respond accordingly.[3]

As before, any stable profile must be a PBE-H. Moreover, stability also imposes additional conditions for profiles π that are on-path strict for the receiver or are such that the sender types' payoffs would not be changed if the receiver deviated.[4] For such a profile to be stable, it must be that, for every signal s where $u_1(\theta, s, a) < u_1(\theta, \pi)$ for all $a \in BR(\overline{\Theta}(s, \pi), s)$ and $\theta \notin \overline{\Theta}(s, \pi)$, there is some $m \in M$ such that $\pi_2(\cdot|s, m) \in \Delta(BR(\overline{\Theta}(s, \pi), s))$. Aside from the qualifying condition $u_1(\theta, s, a) < u_1(\theta, \pi)$ for all $a \in BR(\overline{\Theta}(s, \pi), s)$ and $\theta \notin \overline{\Theta}(s, \pi)$, this requirement is the same as Condition 2 of Definition 3. Combined, these conditions are weaker than JCE, so they are satisfied

[3] The receiver responding to "credible" statements in this way is similar to the motivation underlying "credible robust neologisms" in Clark (2020).

[4] These restrictions on π guarantee that a typical receiver agent will learn the equilibrium payoffs of the sender types with high probability.

by the equilibria we focus on in Examples 3 and 2. The conditions coincide with JCE in Example 4 provided that the game is altered to have sufficiently fine action spaces. Unlike JCE, the conditions are satisfied by the D1 equilibrium in Example 1, but there are other games in which the conditions rule out D1 equilibria.

OA.8.2 Strengthening Initial Trust

Suppose that we strengthen initial trust to require that for any $s \in S$ and $\widetilde{\Theta}, \widetilde{\Theta}' \subseteq \Theta$, if the receiver has never seen a type outside of $\widetilde{\Theta} \cup \widetilde{\Theta}'$ play $(s, m_{s,\widetilde{\Theta}})$, then their response to a first instance of $(s, m_{s,\widetilde{\Theta}})$ will belong to $BR(\widetilde{\Theta} \cup \widetilde{\Theta}', s)$. This means that a receiver who has only observed types in $\widetilde{\Theta}'$ deceitfully play $(s, m_{s,\widetilde{\Theta}})$ puts high probability on the sender type being in either $\widetilde{\Theta}$ or $\widetilde{\Theta}'$ after observing this signal-message pair. This seems plausible; however, we focus on initial trust because of JCE is simpler and easier to apply than its iterated version.

The stable profiles then satisfy an iterated version of JCE, which itself is stronger than the *Iterated Intuitive Criterion* (Cho and Kreps, 1987) and *co-divinity* (Sobel, Stole and Zapater, 1990). Moreover, it is not nested with NWBR, but it is weaker than the refinement obtained by iteratively applying NWBR.

Fix $s \in S$ and $\pi \in \Pi_1 \times \Pi_2$. Consider the following iterated version of the JCE procedure for computing the set of justified types. Initialize $\overline{\Theta}^0(s, \pi) = \overline{\Theta}(s, \pi)$. For $n \in \{1, 2, 3, ...\}$, let

$$\widetilde{D}^n_\theta(s, \pi) = \{\alpha \in \Delta(BR(\overline{\Theta}^{n-1}(s, \pi), s)) : u_1(\theta, s, \alpha) > u_1(\theta, \pi)\},$$

$$\widetilde{D}^{0,n}_\theta(s, \pi) = \{\alpha \in \Delta(BR(\overline{\Theta}^{n-1}(s, \pi), s)) : u_1(\theta, s, \alpha) = u_1(\theta, \pi)\},$$

$$\Theta^{\dagger,n}(s, \pi) = \{\theta \in \Theta : \widetilde{D}^n_\theta(s, \pi) \cup \widetilde{D}^{0,n}_\theta(s, \pi) \not\subseteq \cup_{\theta' \neq \theta} \widetilde{D}_{\theta'}(s, \pi)\},$$

$$\overline{\Theta}^n(s, \pi) = \begin{cases} \Theta^{\dagger,n}(s, \pi) & \text{if } \Theta^{\dagger,n}(s, \pi) \neq \emptyset \\ \overline{\Theta}^{n-1}(s, \pi) & \text{if } \Theta^{\dagger,n}(s, \pi) = \emptyset \end{cases}.$$

Set $\overline{\Theta}^\infty(s, \pi) = \cap_{n \in \mathbb{N}} \overline{\Theta}^n(s, \pi)$. Note that $\overline{\Theta}^{n+1}(s, \pi) \subseteq \overline{\Theta}^n(s, \pi)$ for all n and that

$\overline{\Theta}^{\infty}(s,\pi) \subseteq \overline{\Theta}^{0}(s,\pi) = \overline{\Theta}(s,\pi)$.

Under this strengthening of initial trust, every stable profile π must satisfy the following requirement: For every signal s, there is some $m \in M$ such that $\pi_2(\cdot|s,m) \in \Delta(BR(\overline{\Theta}^{\infty}(s,\pi),s))$. We refer to PBE-H that satisfy this requirement as *strongly justified communication equilibria*.

The proof proceeds by using similar arguments to the proof of Theorem 1 to inductively establish that $\pi_2(\cdot|s,m_{s,\overline{\Theta}^{\infty}(s,\pi)}) \in \Delta(BR(\overline{\Theta}^{n}(s,\pi),s))$ for all $n \in \mathbb{N}$.

OA.8.3 Costs of Lying

Suppose that we allow the sender's utility function $u_1 : \Theta \times S \times M \times A \to \mathbb{R}$ to depend on the sender's message m in the following way: For all $\theta \in \Theta$ and $\Theta', \Theta'' \subseteq \Theta$ such that $\theta \in \Theta' \cap \Theta''$, and $\Theta''' \subseteq \Theta$ such that $\theta \notin \Theta'''$, $u_1(\theta,s,m_{s,\Theta'},a) = u_1(\theta,s,m_{s,\Theta''},a) \geq u_1(\theta,s,m_{s,\Theta'''},a)$ for all $s \in S$ and $a \in A$. Here lying is weakly costly for the sender in that, for a given s and a, the sender gets a lower payoff from a message that represents a set of types to which they do not belong. For simplicity, we assume that all messages that represent a set containing the true type give the sender the same payoff.

For each signal s, message m, and profile π, we will define a set of types $\overline{\Theta}(s,m,\pi)$ that is analogous to the set of justified types in our main setting where m does not impact payoffs. To do this, first set

$$\widetilde{D}_{\theta}(s,m,\pi) = \{\alpha \in \Delta(BR(\Theta,s)) : u_1(\theta,s,m,\alpha) > u_1(\theta,\pi)\},$$

$$\widetilde{D}^{0}_{\theta}(s,m,\pi) = \{\alpha \in \Delta(BR(\Theta,s)) : u_1(\theta,s,m,\alpha) = u_1(\theta,\pi)\},$$

and

$$\Theta^{\dagger}(s,m,\pi) = \{\theta \in \Theta : \widetilde{D}_{\theta}(s,m,\pi) \cup \widetilde{D}^{0}_{\theta}(s,m,\pi) \not\subseteq \cup_{\theta' \neq \theta} \widetilde{D}_{\theta'}(s,m,\pi)\}$$

Then let

$$\overline{\Theta}(s,m,\pi) = \begin{cases} \Theta^{\dagger}(s,m,\pi) & \text{if } \Theta^{\dagger}(s,m,\pi) \neq \emptyset \\ \Theta & \text{if } \Theta^{\dagger}(s,m,\pi) = \emptyset \end{cases}.$$

Under initial trust, a similar proof to that of Theorem 1 shows that any stable profile π must satisfy the following requirement: $\pi_2(\cdot|s, m_{s,\overline{\Theta}(s,\Theta,\pi)}) \in \Delta(BR(\overline{\Theta}(s,\Theta,\pi), s))$ for all $s \in S$. When the sender's message is payoff irrelevant, $\overline{\Theta}(s,m,\pi) = \overline{\Theta}(s,\pi)$, so this requirement implies Condition 2 of Definition 3. While lying costs make it less appealing for a non-justified type to falsely represent themself as justified, they can change the set of equilibria, so it is hard to give a precise summary of their effect in general games.

OA.9 Stability Under a More General Limit

In this section, we study steady state aggregate play in the more general limit where first γ_1 tends to 1, and then δ and γ_2 tend to 1, without any restrictions on the relative speed with which δ and γ_2 converge. Formally, we consider $\lim_{(\delta,\gamma_2)\to(1,1)} \lim_{\gamma_1\to 1} \Pi^*(g,\delta,\gamma_1,\gamma_2)$. We will call these the *stable* profiles.*

Definition OA 3. *Strategy profile* π *is* **stable*** *if there is a sequence* $\{\delta_j\}_{j\in\mathbb{N}} \to 1$*, sequence* $\{\gamma_{2,j}\}_{j\in\mathbb{N}} \to 1$*, and sequences* $\{\gamma_{1,j,k}\}_{j,k\in\mathbb{N}}$ *with* $\lim_{k\to\infty} \gamma_{1,j,k} = 1$ *for all* j*, such that* $\pi = \lim_{j\to\infty} \lim_{k\to\infty} \pi_{j,k}$ *for some sequence* $\pi_{j,k} \in \Pi^*(g, \delta_{1,j}, \gamma_{1,j,k}, \gamma_{2,j})$*.*

Since every stable profile is also stable*, it follows that stable* profiles exist.

Corollary OA 2. *Stable* strategy profiles exist.*

As with stability, there is a strong relationship between the stable* profiles and the set of JCE.

Definition OA 4. *Strategy profile* π *has* **strong incentives** *if, for every off-path* s *and* $\theta \notin \overline{\Theta}(s,\pi)$*, there is some on-path* (s',m') *such that* $u_1(\theta,s',a) > u_1(\theta,s,\pi_2(\cdot|s,m_{s,\overline{\Theta}(s,\pi)}))$

for all $a \in BR(p_{(s',m')}, s')$, where $p_{(s',m')}$ is the posterior belief given (s',m') obtained from π_1 and Bayes' rule.

A strategy profile has strong incentives if for every off-path s, every type would obtain a strictly lower payoff from playing $(s, m_{s,\overline{\Theta}(s,\pi)})$ than they would from playing some on-path signal-message pair when the receiver responds with any best response to the corresponding posterior.

Theorem OA 1. *Suppose that the density of the prior of the sender agents is everywhere positive. If π is stable* and has strong incentives, then it is a JCE.*

Theorem OA 1 says that a profile with strong incentives can be stable* only if it is a JCE. The assumption of strong incentives is vacuous if all signals are played with positive probability in π. Also, note that $u_1(\theta, \pi) > u_1(\theta, s, m_{s,\overline{\theta}(s,\pi)})$ for an arbitrary signal s and profile π whenever $\theta \notin \overline{\Theta}(s, \pi)$. Thus, every profile that is on-path strict for the receiver has strong incentives.[5]

The remainder of this section is devoted to the proof of Theorem OA 1. The argument that every stable* profile is a PBE-H proceeds very similarly to that for the stable profiles. The following lemma affirms the optimality of the aggregate sender play given the aggregate receiver play.

Lemma OA 1. *Suppose that π is stable*. Then for each $\theta \in \Theta$, $\pi_1(\cdot|\theta)$ puts support only on those sender signal-message pairs that are optimal for type θ under the receiver behavior strategy π_2.*

The next lemma shows that aggregate receiver play is a best response to (on-path) aggregate play by the senders in a stable* profile.

Lemma OA 2. *Suppose that π is stable*. Then for any sender signal-message pair $(s, m) \in S \times M$ that occurs with positive probability under π, $\pi_2(\cdot|s, m)$ puts support only on receiver actions that are best-responses to s and the posterior belief induced by λ and $\{\pi_1(s, m|\theta)\}_{\theta \in \Theta}$.*

[5]Another sufficient condition is that no sender type would be hurt if the receiver were to change their response to some on-path signal-message pair, as is the case when all types choose an "exit" option.

We omit the proofs of Lemma OA 1 and Lemma OA 2, which are quite similar to the proofs of Lemma 1 and Lemma 2, respectively.

Lemma OA 3 below shows that when π is a stable profile that has strong incentives, the aggregate receiver response to any $(s, m_{s,\overline{\Theta}(s,\pi)})$ must be supported on $BR(\overline{\Theta}(s,\pi), s)$.

Lemma OA 3. *Suppose that π is stable* and has strong incentives. Then $\pi_2(\cdot|s, m_{s,\overline{\Theta}(s,\pi)}) \in \Delta(BR(\overline{\Theta}(s,\pi), s))$ for all $s \in S$.*

We prove Lemma OA 3 in the following subsection, but first we use Lemmas OA 1, OA 2, and OA 3 to prove Theorem OA 1.

Proof of Theorem OA 1. Lemma OA 1 implies Condition 1 of the definition of PBE-H, and Lemma OA 2 implies Condition 2. As before, Condition 3 of Definition 1 follows from the fact that the receivers in our model myopically optimize. Finally, the additional condition in Definition 3 follows from Lemma OA 3 and the assumption that π has strong incentives. ■

OA.9.1 Proof of Lemma OA 3

The following lemma relates the receiver's continuation parameter to the probability the aggregate receiver response to any on-path signal-message pair places on the corresponding receiver best responses.

Lemma OA 4. *Fix a strategy profile π. Let X^{on} be the set of sender signal-message pairs that are on-path under π_1, and let $p_{(s,m)}$ be the posterior belief given $(s,m) \in X^{on}$ that is obtained from π_1 and Bayes' rule. There are $\nu, \eta > 0$ such that, for every $\pi_1' \in \Pi_1$ satisfying $\max_{(\theta,s,m)\in\Theta\times S\times M} |\pi_1'(s,m|\theta) - \pi_1(s,m|\theta)| < \nu$ and all $\delta, \gamma_1, \gamma_2 \in [0,1)$,*

$$\mathscr{R}_2^{\gamma_2}(\pi_1')(BR(p_{(s,m)}, s)|(s,m)) \geq 1 - \eta(1-\gamma_2)$$

for all $(s,m) \in X^{on}$.

Proof. Let $q(\theta, s, m) = \lambda(\theta)\pi_1(s, m|\theta)$ be the distribution over sender types, signals, and messages induced by λ and π_1. For $\varepsilon > 0$, let $Q_\varepsilon = \{q' \in \Delta(\Theta \times S \times M) : \max_{(\theta,s,m)\in\Theta\times S\times M} |q'(\theta, s, m) - q(\theta, s, m)| \leq \varepsilon\}$. Because best responses are upper hemi-continous, there exists $\varepsilon > 0$ such that every receiver whose belief $\widetilde{g}_2 \in \Delta(\Delta(\Theta\times S\times M))$ puts probability at least $1 - \varepsilon$ on Q_ε will respond to every $(s, m) \in X^{\text{on}}$ with some action belonging to $BR(p_{(s,m)}, s)$.

Given the non-doctrinaire prior g_2, Theorem 4.2 of Diaconis and Freedman (1990) implies that there is some $T > 0$ such that a receiver who has lived more than T periods assigns posterior probability of at least $1 - \varepsilon$ to probability distributions q' within $\varepsilon/3$ distance (in the sup-norm metric) of the empirical distribution they have observed.

We provide a lower bound on the share of receivers who have lived more than T periods and who have observed an empirical distribution within $\varepsilon/3$ distance of the true distribution $q' \in \Delta(\Theta \times S \times M)$. By Hoeffding's inequality, the probability that the fraction of (θ, s, m) observations is outside of $[q'(\theta, s, m) - \varepsilon/3, q'(\theta, s, m) + \varepsilon/3]$ for a receiver with t observations is less than $2e^{-\frac{2\varepsilon^2}{9}t}$, so the probability that the empirical distribution of a receiver with t observations is greater than $\varepsilon/3$ distance from q' is no more than $2|S||M|e^{-\frac{2\varepsilon^2}{9}t}$. Thus, the share of receivers who have lived longer than T periods and who have observed an empirical distribution within $\varepsilon/3$ distance of q' is at least

$$\begin{aligned}
\sum_{t=T}^{\infty}(1-\gamma_2)\gamma_2^t\left(1 - 2|S||M|e^{-\frac{2\varepsilon^2}{9}t}\right) &= \gamma_2^T - \frac{2|S||M|(1-\gamma_2)\gamma_2^T e^{-\frac{2\varepsilon^2}{9}T}}{1-\gamma_2 e^{-\frac{2\varepsilon^2}{9}}}, \\
&= 1 - \left(\frac{1-\gamma_2^T}{1-\gamma_2} + \frac{2|S||M|\gamma_2^T e^{-\frac{2\varepsilon^2}{9}T}}{1-\gamma_2 e^{-\frac{2\varepsilon^2}{9}}}\right)(1-\gamma_2), \\
&\geq 1 - \left(T + \frac{2|S||M|}{1-e^{-\frac{2\varepsilon^2}{9}}}\right)(1-\gamma_2),
\end{aligned}$$

where the inequality follows from the facts that $(1-\gamma_2^T)/(1-\gamma_2) < T$ and $\gamma_2^T e^{-\frac{2\varepsilon^2 T}{9}}/(1-\gamma_2 e^{-\frac{2\varepsilon^2}{9}}) < 1/(1-e^{-\frac{2\varepsilon^2}{9}})$ for all $\gamma_2 \in [0, 1)$.

Let $\eta = T + 2|S||M|/(1 - e^{-\frac{2\varepsilon^2}{9}})$, and let $\nu > 0$ be such that, for every $\pi_1' \in \Pi_1$ satisfying $\max_{(\theta,s,m)\in\Theta\times S\times M}|\pi_1'(s,m|\theta)-\pi_1(s,m|\theta)|<\nu$, the corresponding distribution over sender types, signals, and messages belongs to $Q_{\varepsilon/3}$. It follows from the arguments above that, for all π_1' within ν distance (in the sup-norm metric) of π_1, the steady-state share of receivers who respond to each $(s,m)\in X^{\text{on}}$ with some element of $BR(p_{(s,m)},s)$ is at least $1-\eta(1-\gamma_2)$. ■

The next lemma builds on Lemma OA 4 to show that, in a sequence of steady states converging to a stable* profile with strong incentives, the ratio of the aggregate probability of a non-justified type playing $(s,m_{s,\overline{\Theta}(s,\pi)})$ to the expected lifetime of a receiver agent approaches 0.

Lemma OA 5. *Fix a stable* strategy profile π with strong incentives. Let $\{\pi_{j,k}\in\Pi^*(g,\delta_j,\gamma_{1,j,k},\gamma_{2,j})\}_{j,k\in\mathbb{N}}$ be a sequence of steady state profiles such that $\lim_{j\to\infty}\lim_{k\to\infty}\pi_{j,k}=\pi$, where $\lim_{j\to\infty}\delta_j=1$, $\lim_{j\to\infty}\gamma_{2,j}=1$, and $\lim_{k\to\infty}\delta_{j,k}=1$ for all j. For every $\varepsilon>0$, there exists some $J\in\mathbb{N}$ and function $K:\mathbb{N}\to\mathbb{N}$ such that*

$$\pi_{1,j,k}(s,m_{s,\overline{\Theta}(s,\pi)}|\theta)\leq\varepsilon(1-\gamma_{2,j})$$

for all s, $\theta\notin\overline{\Theta}(s,\pi)$, $j>J$, and $k>K(j)$.

Proof. By Lemma OA 4 and the fact that $\lim_{j\to\infty}\lim_{k\to\infty}\pi_{j,k}=\pi$, there exists some $\eta>0$, $J'\in\mathbb{N}$, and function $K':\mathbb{N}\to\mathbb{N}$ such that

$$\pi_{2,j,k}(BR(p_{(s,m)},s)|(s,m))\geq 1-\eta(1-\gamma_{2,j}) \tag{1}$$

for all (s,m) on-path under π_1, $j>J'$, and $k>K'(j)$.

Fix a signal s and type θ such that $\theta\notin\overline{\Theta}(s,\pi)$. Since π has strong incentives, there is some (s',m') that is on-path under π_1 such that $u_1(\theta,s',a)>u_1(\theta,s,\pi_2(\cdot|s,m_{s,\overline{\Theta}(s,\pi)}))$ for all $a\in BR(p_{(s',m')},s')$. For any $\alpha\in\Delta(A)$ and $z>0$, let $\mathcal{A}_{(\alpha,z)}=\{\alpha'\in\Delta(A):\max_{a\in A}|\alpha'[a]-\alpha[a]|\leq z\}$ be the set of mixtures over A that are no greater than z

away from α in the sup-norm metric. Let $\nu > 0$ be such that

$$(1-\nu)u_1(\theta, s', a) + \nu \min_{a' \in A} u_1(\theta, s', a') > u_1(\theta, s, \alpha) + \nu \tag{2}$$

for all $a \in BR(p_{(s',m')}, s')$ and $\alpha \in \mathcal{A}_{(\pi_2(\cdot|s, m_{s,\overline{\Theta}(s,\pi)}), \nu)}$.

Suppose that a sender has played (s', m') at least $N > 0$ times. Combining Equation 1 with Lemma A.1 of Fudenberg and Levine (2006) implies that the probability that the fraction of times the sender observed a receiver play something outside of $BR(p_{(s',m')}, s')$ in response to (s', m') exceeds $\nu/2$ is no more than $2^{11}\eta(1-\gamma_{2,j})/(3\nu^4 N)$. For a fixed $\varepsilon > 0$, let $N_{(s',m')}$ be such that $2^{11}\eta/(3\nu^4 N_{(s',m')}) < \varepsilon/4$. For such an $N_{(s',m')}$, it follows that $2^{11}\eta(1-\gamma_{2,j})/(3\nu^4 N_{(s',m')}) < \varepsilon(1-\gamma_{2,j})/4$.

By the assumption that the sender's prior has a density $g_1(\pi_2)$ that is everywhere positive and continuous in $\pi_2 \in \Pi_2$, we can find a lower bound on the probability that certain senders put on the receiver aggregate response to (s', m') playing an element of $BR(p_{(s',m')}, s')$ with probability at least $1-\nu$. In particular, we will show there is a lower bound $\zeta > 0$ on the probability that the aggregate receiver response to (s', m') puts probability at least $1-\nu$ on $BR(p_{(s',m')}, s')$ as determined by two classes of sender agents: (1) a sender agent who has played (s', m') fewer than $N_{(s',m')}$ times, and (2) a sender agent who has played (s', m') more than $N_{(s',m')}$ times and observed a response in $BR(p_{(s',m')}, s')$ greater than a fraction $1-\nu/2$ of the times. From the preceding paragraph, the share of sender agents who fall into either of these two classes exceeds $1-\varepsilon(1-\gamma_{2,j})/4$.

Consider a sender who, for each $a \in A$, has n_a observations of a receiver responding to (s', m') with a. Then such a sender puts probability at least

$$\frac{\min_{\pi_2 \in \Pi_2} g_1(\pi_2) \int_{\{\alpha \in \Delta(A) : \alpha[BR(p_{(s',m')}, s')] \geq 1-\nu\}} \Pi_{a \in A} \alpha[a]^{n_a}}{\max_{\pi_2 \in \Pi_2} g_1(\pi_2) \int_{\Delta(A)} \Pi_{a \in A} \alpha[a]^{n_a}}$$

on the set of aggregate receiver responses to (s', m') that have probability weakly greater than $1-\nu$ on $BR(p_{(s',m')}, s')$. This expression is uniformly bounded away from

0 when there are fewer than $N_{(s',m')}$ observations. Moreover, Theorem 4.2 of Diaconis and Freedman (1990) implies that this expression is uniformly bounded away from 0 when there are more than $N_{(s',m')}$ observations and the fraction of these observations where the receiver responding with some element of $BR(p_{(s',m')}, s')$ exceeds $1 - \nu/2$.

By similar arguments, there is some $N'_s \in \mathbb{N}$ such that, for a sender who has played $(s, m_{s,\overline{\Theta}(s,\pi)})$ at least N'_s times, the sender's expectation of the aggregate receiver response to $(s, m_{s,\overline{\Theta}(s,\pi)})$ is within $\nu/3$ (in the sup-norm metric) of the empirical response the sender has observed. Moreover, by the law of large numbers, for any $j \in \mathbb{N}$, we can choose some $N'_{s,j} > N'_s$ to be such that there is a probability no greater than $\varepsilon(1 - \gamma_{2,j})/4$ that the empirical response to $(s, m_{s,\overline{\Theta}(s,\pi)})$ observed by a sender who has played $(s, m_{s,\overline{\Theta}(s,\pi)})$ at least $N'_{s,j}$ times is more than $\nu/3$ away from the aggregate receiver response $\pi_{2,j,k}(\cdot|s, m_{s,\overline{\Theta}(s,\pi)})$. Let $J'' \in \mathbb{N}$ and $K'' : \mathbb{N} \to \mathbb{N}$ be such that $\max_{a \in A} |\pi_{2,j,k}(a|s, m_{s,\overline{\Theta}(s,\pi)}) - \pi_2(a|s, m_{s,\overline{\Theta}(s,\pi)})| < \nu/3$ for all $j > J''$ and $k > K''(j)$. It follows that, for all such j and k, the probability that $\mathcal{A}_{(\pi_2(\cdot|s,m_{s,\overline{\Theta}(s,\pi)}),\nu)}$ contains the expectation of the aggregate receiver response to $(s, m_{s,\overline{\Theta}(s,\pi)})$, as evaluated by a sender who has played $(s, m_{s,\overline{\Theta}(s,\pi)})$ at least $N'_{s,j}$ times, exceeds $1 - \varepsilon(1 - \gamma_{2,j})/4$.

Consider a sender belief $\widetilde{g}_1 \in \Delta(\Pi_2)$ that satisfies

$$\begin{aligned} &\widetilde{g}_1(\pi_2(BR(p_{(s',m')}, s')|s', m') \geq 1 - \nu) \geq \zeta, \\ &\widetilde{g}_1(\pi_2(\cdot|s, m_{s,\overline{\theta}(s,\pi)}) \in \mathcal{A}_{(\pi_2(\cdot|s,m_{s,\overline{\Theta}(s,\pi)}),\nu)}) \geq 1 - \frac{1}{2}\zeta. \end{aligned} \tag{3}$$

The first inequality says that the belief puts probability at least ζ on aggregate receiver responses to (s', m') that play an element of $BR(p_{(s',m')}, s')$ with probability weakly greater than $1 - \nu$. The second inequality says that the belief puts probability at least $1-\zeta/2$ on the aggregate receiver response to $(s, m_{s,\overline{\Theta}(s,\pi)})$ belonging to $\mathcal{A}_{(\pi_2(\cdot|s,m_{s,\overline{\Theta}(s,\pi)}),\nu)}$. By Equation 2, all beliefs satisfying the conditions in (3) must put probability at least $\zeta/2$ on aggregate receiver behavior strategies where playing (s', m') gives a type θ sender an expected payoff at least ν greater than that from playing $(s, m_{s,\overline{\Theta}(s,\pi)})$.

For a type θ sender with any belief that satisfies (3), the expected total lifetime payoff from the optimal policy exceeds the expected total lifetime payoff from only playing $(s, m_{s,\overline{\Theta}(s,\pi)})$ by an amount bounded away from 0 when δ and γ_1 are sufficiently high. In particular, for δ and γ_1 sufficiently close to 1, the difference in the expected payoff from the optimal policy and that from repeatedly playing $(s, m_{s,\overline{\Theta}(s,\pi)})$ exceeds $c = \zeta\nu/4 > 0$. Let $J''' \in \mathbb{N}$ and $K''' : \mathbb{N} \to \mathbb{N}$ be such that, whenever $j > J'''$ and $k > K'''(j)$, δ_j and $\gamma_{1,j,k}$ are sufficiently close to 1 so that this gap in the expected payoffs holds. Then, the version of Corollary 5.5 of Fudenberg and Levine (1993) presented in Fudenberg and He (2018) implies that, for every $j > J'''$, there is some $N''_{s,j}$ such that the share of type θ sender agents who have a belief satisfying the conditions in (3), have played $(s, m_{s,\overline{\Theta}(s,\pi)})$ more than $N''_{s,j}$ times, and are set to play $(s, m_{s,\overline{\Theta}(s,\pi)})$ in the current period is less than $\varepsilon(1 - \gamma_{2,j})/4$ for all $k > K'''(j)$.

Let $J = \max\{J', J'', J'''\}$, $K(j) = \max\{K'(j), K''(j), K'''(j)\}$ for all $j > J$, and $N_{s,j} = \max\{N'_{s,j}, N''_{s,j}\}$ for all $j > J$. Combining the preceding results shows that, when $j > J$ and $k > K(j)$, the share of type θ sender agents who have played $(s, m_{s,\overline{\Theta}(s,\pi)})$ more than $N_{s,j}$ times and are set to play $(s, m_{s,\overline{\Theta}(s,\pi)})$ in the current period is no more than $3\varepsilon(1 - \gamma_{2,j})/4$. Additionally, using the version of Lemma 5.7 of Fudenberg and Levine (1993) presented in Fudenberg and He (2018), it follows that, for all $j > J$, $K(j)$ can also be chosen so that $\pi_{1,j,k}(s, m_{s,\overline{\Theta}(s,\pi)}|\theta)$ exceeds the share of type θ sender agents who have played $(s, m_{s,\overline{\Theta}(s,\pi)})$ more than $N''_{s,j}$ times and are set to play $(s, m_{s,\overline{\Theta}(s,\pi)})$ in the current period by no more than $\varepsilon(1 - \gamma_{2,j})/4$ when $k > K(j)$. Thus, we conclude that $\pi_{1,j,k}(s, m_{s,\overline{\Theta}(s,\pi)}|\theta) \leq \varepsilon(1 - \gamma_2)$ for all $j > J$ and $k > K(j)$. ■

The proof of Lemma OA 3 uses Lemma OA 5 to show that, in a sequence of steady states converging to a stable* profile with strong incentives, the probability that a receiver encounters a non-justified sender type playing some $(s, m_{s,\overline{\Theta}(s,\pi)})$ over the course of their lifetime converges to 0. Initial trust then ensures that the aggregate receiver response to each $(s, m_{s,\overline{\Theta}(s,\pi)})$ is justified.

Proof of Lemma OA 3. Let $\{\pi_{j,k} \in \Pi^*(g, \delta_j, \gamma_{1,j,k}, \gamma_{2,j})\}_{j,k\in\mathbb{N}}$ be a sequence of steady state profiles such that $\lim_{j\to\infty} \lim_{k\to\infty} \pi_{j,k} = \pi$, where $\lim_{j\to\infty} \delta_j = 1$, $\lim_{j\to\infty} \gamma_{2,j} = 1$, and $\lim_{k\to\infty} \delta_{j,k} = 1$ for all j. By Lemma OA 5, for any $\varepsilon > 0$, there exists some $J \in \mathbb{N}$ and some function $K : \mathbb{N} \to \mathbb{N}$ such that $\pi_{1,j,k}(s, m_{s,\overline{\Theta}(s,\pi)}|\theta) \leq \varepsilon(1-\gamma_2)/\lambda(\theta)$ for all $\theta \notin \overline{\Theta}(s,\pi)$, $j > J$, and $k > K(j)$. Thus, when $j > J$ and $k > K(j)$, the probability that a receiver agent in a given period encounters a sender type outside of $\overline{\Theta}(s,\pi)$ playing $(s, m_{s,\overline{\Theta}(s,\pi)})$ is no greater than $\varepsilon(1-\gamma_{2,j})$. It follows that, when $j > J$ and $k > K(j)$, the probability that a receiver agent never encounters a sender type outside of $\overline{\Theta}(s,\pi)$ playing $(s, m_{s,\overline{\Theta}(s,\pi)})$ over the course of their lifetime is at least

$$\sum_{t=0}^{\infty}(1-\gamma_{2,j})\gamma_{2,j}^{t}(1-\varepsilon(1-\gamma_{2,j}))^t = \frac{1}{1+\gamma_{2,j}\varepsilon}.$$

Receivers who have never observed the signal-message pair $(s, m_{s,\overline{\Theta}(s,\pi)})$ played by a type outside of $\overline{\Theta}(s,\pi)$ would respond to this pair with an action belonging to $BR(\overline{\Theta}(s,\pi), s)$. Thus,

$$\pi_2(BR(\overline{\Theta}(s,\pi),s)|s, m_{s,\overline{\Theta}(s,\pi)}) = \lim_{j\to\infty}\lim_{k\to\infty} \pi_{2,j,k}(BR(\overline{\Theta}(s,\pi),s)|s, m_{s,\overline{\Theta}(s,\pi)}) \geq 1/(1+\varepsilon).$$

Since this holds for all $\varepsilon > 0$, we have that $\pi_2(BR(\overline{\Theta}(s,\pi),s)|s, m_{s,\overline{\Theta}(s,\pi)}) = 1$. ∎

OA.10 Details of Alternate Model

Consider a steady-state population of receivers who have geometric lifetimes with continuation probability γ, and are matched with a sender each period with i.i.d. probability p. We show that, when the receivers have expected lifespan $T = 1/(1-\gamma)$ and are expected to have $N_2 = pT$ matches over the course of their lifetime, the distribution of match experience in the receiver population is geometric with hit probability $\tilde{\gamma}_2 = (1-1/T)N_2/(1+(1-1/T)N_2)$. Because the aggregate play of receivers only depends on their experience, it follows that for every steady state in our main learning model given parameters γ_1, δ, and γ_2, there is a steady state in this alternate model given parameters $\gamma = \gamma_1$, δ, and $\tilde{\gamma}_2$ with the same aggregate strategy profile.

Lemma OA 6. *If receivers have geometric lifetimes with expected lifespan T and are expected to have N_2 matches over the course of their lifetime, then the steady-state share of receivers who have previously been matched $n \in \mathbb{N}$ times is $(1-\tilde{\gamma}_2)\tilde{\gamma}_2^n$, where*

$$\tilde{\gamma}_2 = \frac{\left(1 - \frac{1}{T}\right) N_2}{1 + \left(1 - \frac{1}{T}\right) N_2}.$$

Proof. Denote the steady-state share of receivers who have previously had n matches by $\tilde{\mu}_2[n]$. We first derive $\tilde{\mu}_2[0]$. Since $1-\gamma$ is the share of newborn receivers and $\gamma(1-p)\tilde{\mu}_2[0]$ is the share of non-newborn receivers who have never been matched, it follows that $\tilde{\mu}_2[0] = (1-\gamma) + \gamma(1-p)\tilde{\mu}_2[0]$. Solving this gives

$$\tilde{\mu}_2[0] = \frac{1-\gamma}{1-\gamma+\gamma p}. \tag{OA 1}$$

Now we derive a recursive expression relating $\tilde{\mu}_2[n]$ to $\tilde{\mu}_2[n-1]$ for $n > 0$. Since $\gamma p \tilde{\mu}_2[n-1]$ is the share of receivers who in the previous period were matched for the nth time and $\gamma(1-p)\tilde{\mu}_2[0]$ is the share of receivers who have been matched n times but were unmatched in the previous period, it follows that $\tilde{\mu}_2[n] = \gamma p \tilde{\mu}_2[n-1] + \gamma(1-p)\tilde{\mu}_2[n]$. Solving this gives

$$\tilde{\mu}_2[n] = \frac{\gamma p}{1-\gamma+\gamma p}\tilde{\mu}_2[n-1]. \tag{OA 2}$$

Combining Equations OA 1 and OA 2 gives

$$\tilde{\mu}_2[n] = \left(1 - \frac{\gamma p}{1-\gamma+\gamma p}\right)\left(\frac{\gamma p}{1-\gamma+\gamma p}\right)^n.$$

Substituting $\gamma = 1 - 1/T$ and $p = N/T$ renders

$$\tilde{\mu}_2[n] = \left(1 - \frac{\left(1-\frac{1}{T}\right)N_2}{1+\left(1-\frac{1}{T}\right)N_2}\right)\left(\frac{\left(1-\frac{1}{T}\right)N_2}{1+\left(1-\frac{1}{T}\right)N_2}\right)^n$$

as desired. ■

References

Camerer, Colin, and Ari Vepsalainen. 1988. "The Economic Efficiency of Corporate Culture." *Strategic Management Journal*, 9: 115–126.

Cho, In-Koo, and David M. Kreps. 1987. "Signaling Games and Stable Equilibria." *Quarterly Journal of Economics*, 102(2): 179–221.

Clark, Daniel. 2020. "Robust Neologism Proofness in Signaling Games." `http://economics.mit.edu/files/16758`.

Diaconis, Persi, and David Freedman. 1990. "On the Uniform Consistency of Bayes Estimates for Multinomial Probabilities." *The Annals of Statistics*, 18(3): 1317–1327.

Fudenberg, Drew, and David K. Levine. 1993. "Steady State Learning and Nash Equilibrium." *Econometrica*, 61(3): 547–573.

Fudenberg, Drew, and David K. Levine. 2006. "Superstition and Rational Learning." *American Economic Review*, 96(3): 630–651.

Fudenberg, Drew, and Kevin He. 2018. "Learning and Type Compatibility in Signaling Games." *Econometrica*, 86(4): 1215–1255.

Fudenberg, Drew, and Kevin He. 2020. "Payoff Information and Learning in Signaling Games." *Games and Economic Behavior*, 120: 96–120.

Kreps, David M. 1990. "Perspectives on Positive Political Economy." Chapter Corporate Culture and Economic Theory, 90–142. Cambridge University Press.

Sobel, Joel, Lars Stole, and Iñigo Zapater. 1990. "Fixed-Equilibrium Rationalizability in Signaling Games." *Journal of Economic Theory*, 52(2): 304–331.