Online Appendix for "Imperfect Public Monitoring with a Fear of Signal Distortion"

Vivek Bhattacharya[1], Lucas Manuelli[2], and Ludwig Straub[3]

[1]Department of Economics, Northwestern University
[2]CSAIL, Massachusetts Institute of Technology
[3]Department of Economics, Massachusetts Institute of Technology

December 1, 2017

B. Incorporating Ambiguity Aversion

In this appendix, we begin with a standard equilibrium concept to solve the game presented in Section 2.1. We then add ambiguity aversion to this setup and show that it leads to the concept of distortion equilibrium as we defined it in Section 2.2.

B.1. Sequential Equilibrium

In this subsection, we will consider a sequential equilibrium in this game (Kreps and Wilson (1982)), which consists of a set of mixed strategies at each information set along with beliefs over nodes at each set such that actions are sequentially rational, beliefs are formed by Bayes' rule whenever it applies, and beliefs are consistent (in the sense defined by Kreps and Wilson) for information sets that have probability zero in equilibrium. Denote by $p_i(h_i) \in [0,1]$ the probability that player i accepts μ over $\hat{y}$ at information set h_i, and let $\nu_i(a_{-i}, O|h_i)$ be i's belief that his opponent played a_{-i} and the order is O, conditional on being at information set h_i. To simplify notation, consider only player 1 and let $\nu(1|h)$ be his belief that he is picking first, at information set h. The expected utility to player 1 from accepting the alternate distribution μ is

$$U_1(\mu|h) = \nu(1|h)\chi(\mu, h) + (1 - \nu(1|h))\mathbb{E}_{\mu[y]} w_1(y), \tag{O.1}$$

where $\chi(1, h)$ is the expected profit of player 1 from choosing μ at information set h, conditional on being first to alter the signal. We can write

$$\chi(\mu, h) = \sum_{a_2 \in A_2} \frac{\nu(a_2, 1|h)}{\nu(1|h)} \mathbb{E}_{\mu[y']} \mathbb{E}_{F(a_2,y')[\mu']} \left\{ w_1(\mu') p_2(a_2, y', \mu') + w_1(y') \left(1 - p_2(a_2, y', \mu')\right)\right\}, \tag{O.2}$$

where $w_1(\mu')$ is shorthand for $\mathbb{E}_{\mu'[y'']} w_1(y'')$. Of course, the expected utility to player 1 from accepting the temporary signal $\hat{y}$ is simply $U_1(\delta_{\hat{y}}|h)$ from (O.1), where $\delta_{\hat{y}}$ is a distribution that places probability 1 on $\hat{y}$.

Contact. vivek.bhattacharya@northwestern.edu, manuelli@mit.edu, straub@mit.edu.

Sequential rationality for player 1 at these information sets, therefore, requires that

$$p_1(a_1, y', \mu) = \begin{cases} 1 & \text{if } U_1(\mu|h) > U_1(\delta_{\hat{y}}|h) \\ [0,1] & \text{if } U_1(\mu|h) = U_1(\delta_{\hat{y}}|h) \\ 0 & \text{if } U_1(\mu|h) < U_1(\delta_{\hat{y}}|h) \end{cases}. \tag{O.3}$$

Bayesian updating requires that

$$\nu(a_2, 1|a_1, \hat{y}, \mu) = \pi(a_1, a_2)(\hat{y}) \cdot dF(\hat{y})(\mu) \cdot \sigma_1(a_1)\sigma_2(a_2)\gamma \bigg/ \Bigg[\sum_{a_2' \in A_2} \pi(a_1, a_2')(\hat{y}) \cdot dF(\hat{y})(\mu) \cdot \sigma_1(a_1)\sigma_2(a_2')\gamma$$
$$+ \int_{\mu'} \sum_{a_2' \in A_2} \pi(a_1, a_2')(\hat{y}) \cdot dF(\hat{y})(\mu') \cdot (1 - p_2(a_2, \hat{y}, \mu')) \cdot dF(a_2, \hat{y})(\mu) \cdot \sigma_1(a_1)\sigma_2(a_2')(1-\gamma)$$
$$+ \int_{\mu'} \sum_{y' \in Y} \sum_{a_2' \in A_2} \pi(a_1, a_2')(y') \cdot dF(y')(\mu') \cdot p_2(a_2, y', \mu') \cdot \mu'(\hat{y}) \cdot dF(a_2, \hat{y})(\mu) \cdot \sigma_1(a_1)\sigma_2(a_2')(1-\gamma) \Bigg]. \tag{O.4}$$

This expression is extensive, but it captures that the game can arrive at the information set $(a_1, \hat{y}, \mu)$ in three ways: (i) player 1 could be first and have been offered the temporary signal $\hat{y}$ and the alternate μ, (ii) player 1 could be second, and $\hat{y}$ may be the signal because player 2 refused to change it, and (iii) player 1 could be second, and $\hat{y}$ may be the signal because it was the draw from player 1's alternate signal. The probability of reaching this node is zero if $\sigma_1(a_1) = 0$ or if π and the F's are such that either $\hat{y}$ or μ has probability zero following the action a_1. If the probability of this node is zero due to $\sigma_1(a_1) = 0$, we follow Kreps and Wilson and define $\nu(a_2, 1|a_1, \hat{y}, \mu)$ using consistency. If certain nodes are not reached due to π or the F's, then we can simply ignore them in constructing the equilibrium.

Finally, the mixed strategies $\sigma_i \in \Delta A_i$ chosen in the first stage of the game (i.e., when playing G) must by optimal given the future path of play (i.e., p_1 and p_2 defined in (O.3)). and given the opponents' strategy $\sigma_{-i} \in \Delta A_{-i}$. Together with the conditions (O.3), (O.4), and the consistency requirement from Kreps and Wilson (1982), we have a sequential equilibrium of the game.

B.2. Adding Ambiguity Aversion

We now add ambiguity-aversion to the payoffs presented in the previous subsection. Throughout this derivation, restrict attention to signal structures π such that $\pi(a)$ has full support on Y; we will briefly discuss relaxing this restriction at the end of the subsection. As mentioned in Section 2.2, we assume that agents only know that $\gamma \in [0,1]$, and that agents believe each F is an element of $\Delta(\Delta(Y))^0$, or the set of distributions over ΔY that have full support.

Such an agent, therefore, would evaluate payoffs according to

$$\tilde{U}_1(\mu|h) = \inf_{\substack{\gamma \in [0,1] \\ F \in \Delta(\Delta(Y))^0}} \nu(1|h)\chi(\mu, h) + (1 - \nu(1|h))\mathbb{E}_{\mu[y]} w_1(y), \tag{O.5}$$

where $\nu(1|h)$ and $\chi(1|h)$ are defined as in (O.2) and (O.4), and $p_2(\cdot)$ is taken as fixed. By F in the infimum we refer to all three different distributions, $F, F(a_1,\cdot), F(a_2,\cdot)$ over alternate signal distributions. Note that we have to take the infimum instead of the minimum since $\Delta(\Delta(Y))^0$ is not a closed set. However, because F is an element in $\Delta(\Delta(Y))^0$ and we are currently restricting attention to π such that $\pi(a)$ has full support on Y, we can ensure that all combinations of $(\hat{y},\mu)$ have positive probability following every action profile and every possible ordering.

We are searching for an equilibrium in which (i) p_i is determined by sequential rationality (as in (O.3), with U replaced by $\tilde{U}$), (ii) p_i is either 0 or 1, that is, for simplicity we restrict our attention to pure strategies in the distortion part of the game, and finally (iii) the actions in the initial game G are optimal given the value of $\tilde{U}_i$. In the remainder of this subsection, we show how these assumptions motivate the concept of distortion equilibrium (Section 2.2).

Note than, in equilibrium, agent 2 will weakly prefer to distort the signal in favor of an alternate signal distribution at a history $h_2 = (a_2, y', \mu')$ if $\tilde{U}_2(\mu'|h_2) \geq U_2(\delta_{y'}|h_2)$. Let us denote the set

$$D_2(y') = \{\mu' \in \Delta(Y) : \tilde{U}_2(\mu'|h_2) \geq U_2(\delta_{y'}|h_2)\}, \tag{O.6}$$

as the set of all μ' that player 2 will weakly prefer to $\delta_{y'}$. Note that this set is independent of a_2. Finally, we restrict ourselves to a particular tie-breaking rule for all alternate signal distributions $\mu' \neq \delta_{y'}$: we assume that both agents, when indifferent between distorting and not distorting, choose to distort the signal.

To obtain an explicit characterization of the equilibrium, we begin by computing the value of $\tilde{U}_1(\mu|h)$ in (O.5). This is a convex combination of the set of terms

$$\left\{\left\{\mathbb{E}_{\mu[y']}\mathbb{E}_{F(a_2,y')[\mu']}\left[w_1(\mu')p_2(a_2,y',\mu') + w_1(y')\left(1 - p_2(a_2,y',\mu')\right)\right]\right\}_{a_2}, \mathbb{E}_{\mu[y]}w_1(y)\right\}. \tag{O.7}$$

We have that

$$\min_{\mu' \in D_2(y')} w_1(\mu') \leq w_1(y') \tag{O.8}$$

since by construction $\delta_{y'} \in D_2(y')$. Thus, we can minimize the elements of the set in (O.7) term-by-term by setting

$$F(a_2, y') \equiv \beta\left\{\operatorname*{arg\,min}_{\mu' \in D_2(y')} w_1(\mu')\right\}, \tag{O.9}$$

where $\beta\{\mu\}$ is the distribution that places full mass on μ.

We now regard the global infimum in (O.5). From (O.9) and the tie-breaking rule that $\{\mu' : p_2(a_2, y', \mu') = 1\}$ is identically $D_2(y')$, we have that the global infimum (over all allowed F and γ) in (O.5) is no less than

$$\min\left\{\mathbb{E}_{\mu[y']}\min_{\mu' \in D_2(y')} w_1(\mu'),\quad \mathbb{E}_{\mu[y]}w_1(y)\right\} = \mathbb{E}_{\mu[y']}\min_{\mu' \in D_2(y')} w_1(\mu'), \tag{O.10}$$

where the equality follows from (O.8) again. It then suffices to argue that we can find γ such that, together with the choice of $F(a_2, y')$ in (O.9), achieves this lower bound. It is easy to see from (O.4) that setting $\gamma = 1$ achieves this bound as long as $\pi(a)$ has full support for all a, and (O.10) gives an expression for $\tilde{U}_1(\mu|h)$ from (O.5),

$$\tilde{U}_1(\mu|h) = \mathbb{E}_{\mu[y']} \min_{\mu' \in D_2(y')} w_1(\mu'). \tag{O.11}$$

We can apply similar logic to compute the payoff from a (possibly mixed) action α_1 in the initial game G. This payoff is given by

$$\begin{aligned} V_1(\alpha_1, \alpha_2) = \inf_{F(\hat{y}) \in \Delta(\Delta Y)^0} &(1-\beta) g_1(\alpha) \\ &+ \beta \cdot \mathbb{E}_{\pi(\alpha)[\hat{y}]} \left[\mathbb{E}_{F(\hat{y})[\mu]} \max\{\tilde{U}_1(\mu|(\alpha_i, \hat{y}, \mu)), \tilde{U}_1(\delta_{\hat{y}}|(\alpha_i, \hat{y}, \mu))\} \right], \end{aligned}$$

where we already used the fact that agent 1 believes to go first ($\gamma = 1$).[1] Similar computations show that this expression is minimized when $F(\hat{y})$ is such that the second element in the max exceeds the first. In light of (O.11) this lets us write

$$V_1(\alpha_1, \alpha_2) = (1-\beta) g_1(\alpha) + \beta \cdot \mathbb{E}_{\pi(\alpha)[\hat{y}]} \min_{\mu' \in D_2(\hat{y})} w_1(\mu'). \tag{O.12}$$

All the expressions are symmetric for player 2.

We have thus far restricted attention to signal structures π with full support. The role of this restriction (along with the assumption that $F(a_2, y) \in \Delta(\Delta Y)^0$) is to allow player 1 to hold the belief that he is going first with probability 1 to be rationalizable. Without a full support assumption, it may be possible for player 1 to be certain that he is going second if he is offered a signal $\hat{y}$ that has probability zero given his action a_1 in the game G—and Bayesian updating beliefs about the order of distortion will be ill-defined if $\gamma = 1$. We can instead derive a distortion equilibrium for games where the signal structure π does not have full support as a limit of distortion equilibria for games with the same payoffs but signal structures π'_n that do have full support and converge to π. Such an equilibrium would be identical to the one presented in Section 2.2, We omit the details, but the rationale is that the only role of π in the expressions above is in (O.12), and that expression is continuous in π. In this paper we apply the definition in that section to games with signal structures that do not satisfy this full support assumption.

C. Details of the Extension to an Infinite Horizon

This section contains details of the extension to the infinite horizon model. We first introduce an alternative equilibrium concept that directly mimicks PPE, which we call *PPED*, and show its equivalence with the RDE introduced in Section 7. Second, we prove that PPED payoffs can be characterized recursively as in APS. Third, we prove Theorem 4 in similar fashion as FL. Fourth, we introduce the concept of a *linear PPE,* or *PPEL*, which is a PPE where continuation values line on a positively sloped line after each public history; and we show that PPED and PPEL are essentially equivalent. Finally, we prove our "anti-folk theorem", Theorem 5.

[1] The notation $\tilde{U}_1(\mu|(\alpha_i, \hat{y}, \mu))$ is shorthand for the expectation $\mathbb{E}_{\alpha_i[a_i]} \tilde{U}_1(\mu|(a_i, \hat{y}, \mu))$.

C.1. PPED and RDE

For the purposes of comparison to FL, a more relevant solution concept would be one in which strategies bear direct resemblance to strategies in FL and FLM. We now formulate such a concept and show that it is indeed equivalent to recursive distortion equilibria. As in FL, we focus on public strategies.

Definition O.1 (Public Strategy)**.** *A public strategy for player i is a map $h_i : \mathcal{Y} \to \Delta A_i \times \mathcal{D}$. That is, a public strategy specifies a (possibly mixed) action $\alpha_i(y^{t-1})$ and a distortion $D_i(y^{t-1})(\cdot) \in \mathcal{D}$ for each public history y^{t-1}.*

Now we want to specify payoffs in this infinitely repeated game directly from public strategies. As illustrated by our one period example, and the recursive definition above, the "worst case" for player i occurs when player $-i$ modifies the signal *after* him. Using this insight, we can define continuation payoffs: we will prove later that the continuation payoffs so defined are compatible with the continuation payoffs defined as part of the recursive distortion equilibrium. Fix a public strategy profile as above and suppose we enter period $t+1$ with public history y^t. Define

$$\mathcal{C}(y^t) \equiv \left\{ \left(\mu_1(y^{t+s})(y), \mu_2(y^{t+s})(y)\right)_{s\geq 0} : \mu_i(y^{t+s})(y) \in D_i(y^{t+s})(y) \right\}.$$

That is, $\mathcal{C}(y^t)$ specifies a realization of an alternate distribution μ for each public history y^{t+s} and each temporary signal y. Note that if we specify that the distortion is always such that player 1 distorts first and then player 2 does, then an element $c \in \mathcal{C}$ induces a probability distribution over public histories y^s via the sequence

$$y^t \underset{\pi(\alpha(y^t))}{\longrightarrow} \hat{y} \underset{\mu_1(y^t)(\hat{y})}{\longrightarrow} \hat{\hat{y}} \underset{\mu_2(y^t)(\hat{\hat{y}})}{\longrightarrow} y \underset{\text{current public history}}{\longrightarrow} (y^t, y) \equiv y^{t+1},$$

where each arrow represents a draw from the distribution underneath the arrow. Let $\sigma_1(y^t) : \mathcal{C}(y^t) \to \Delta\left(\left\{y^{t+s}\right\}_{s\geq 0}\right)$ be this map. Then let $\mathcal{S}_1(y^t) \equiv \operatorname{Im}\sigma_1(y^t)$ be its image. Then, the perceived payoff for player 1 of the public strategy profile, starting from history y^t, is

$$v_1(y^t) \equiv \min_{\nu\in\mathcal{S}_1(y^t)} \left\{ (1-\beta)\cdot g_1(\alpha(y^t)) + (1-\beta)\cdot \mathbb{E}_\nu\left[\sum_{s=1}^{\infty} \beta^s g_1(\alpha(y^{t+s}))\right]\right\}. \tag{O.13}$$

For player 2, simply interchange the roles of players 1 and 2 in the above discussion.[2] To preserve the notation from the definition of RDE, given these payoffs $v(y^t)$ from histories $y^t \equiv (y^{t-1}, y_t)$ onward, define continuation payoffs at y^{t-1} as $w(y^{t-1}, y_t) \equiv v(y^t)$.[3] Now, define $\tilde{w}(y^t, \cdot)$ as in (i) in

[2]While the continuation payoffs defined as above are technically treated as definitions of the preferences given public strategies, we should note that they are reasonable given the extensive form game defined in Section 2. There, we showed that when evaluating payoffs, players fear that γ is such that their opponent always modifies the signal after they
2 modifies the signal after player 1 does in the minimization problem. Moreover, note that it is without loss of generality to only consider $\mu \in D_i$ as part of $\mathcal{C}$ since $\delta_y \in D_i(y^{t-1})(y)$ for all public histories; as such, the possibility of the opponent rejecting the proposed signal change is built into the set.

[3]The distinction between the $v(y^{t-1})$ and $w(y^{t-1}, \cdot)$ is pedantic, but it serves to highlight the fact that in an RDE, we

the definition of recursive distortion equilibrium (Definition 7). We have that the natural consistency condition is that

$$D_i(y^{t-1})(y) \equiv \left\{\mu \in \Delta Y : \mathbb{E}_{\mu[y']}\left[\tilde{w}_i(y^{t-1}, y')\right] \geq \tilde{w}_i(y^{t-1}, y)\right\}. \tag{O.14}$$

With this notion, we can define an equilibrium.

Definition O.2 (PPED)**.** *A public perfect equilibrium with distortion (PPED) is a public strategy profile (α, D) such that starting at each public history y^{t-1},*

(i) α_i is a best response to α_{-i} and D_{-i} in the subgame following y^{t-1}, using utility $v_i(y^{t-1})$ defined in (O.13), and

(ii) $D(y^{t-1})(\cdot)$ is consistent with the strategy profile as in (O.14).

The connection between a PPED and an RDE is relatively intuitive, as the best response and consistency conditions in the definition of PPED are exactly as in an RDE. It remains to show that the minimization problem posited in (O.13) is compatible with the payoffs defined as part of the RDE. The following result shows that it is.

Theorem O.1. *Given an RDE $\left(\alpha(y^{t-1}), D(y^{t-1}), w(y^{t-1})\right)$, the public strategy profile $\left(\alpha(y^{t-1}), D(y^{t-1})\right)$ is a PPED. Conversely, if $\left(\alpha(y^{t-1}), D(y^{t-1})\right)$ is a PPED, then $w(y^{t-1}, y_t) \equiv v(y^{t-1}, y_t)$, defined in (O.13), satisfies the Bellman equation.*

Proof. This proof has two pieces. In the first step, we show that an RDE is a PPED. Let the RDE be denoted $\left(\alpha(y^{t-1}), D(y^{t-1}), w(y^{t-1})\right)$. Define $\mathcal{C}'(y^t)$, $\mathcal{S}_i'(y^t)$, and $\sigma_i'(y^t)$ as the one-period analogues of $\mathcal{C}$, $\mathcal{S}_i$, and σ_i from Section 7.[4] We claim that w satisfies the difference equation

$$v_i(y^{t-1}) \equiv w_i(y^{t-2}, y_{t-1}) = \min_{\nu \in \mathcal{S}_i'(y^{t-1})} (1-\beta) g_i(\alpha(y^{t-1})) + \beta \mathbb{E}_{\nu[y]}[w_i(y^{t-1}, y)]. \tag{O.15}$$

It suffices to show that

$$\min_{\nu \in \mathcal{S}_i'(y^{t-1})} \mathbb{E}_{\nu[y]}[w_i(y^{t-1}, y)] = \mathbb{E}_{\pi(\alpha(y^{t-1}))[y]}[\tilde{w}_i(y^{t-1}, y)].$$

view continuation payoffs as vectors (for each player) at each node, and in a PPED, we view payoffs as a single value (for each player) at each node. These notions can of course be interchanged freely.

[4] That is, only consider choices μ_i from $D_i(\cdot)(\cdot)$ for a single period, and the induced distribution is only over the next period's signal.

Note that by consistency and the definition of $\tilde{w}$, we have

$$
\begin{aligned}
&\mathbb{E}_{\pi(\alpha(y^{t-1}))[y_t]}[\tilde{w}_i(y^{t-1}, y_t)] \\
&= \mathbb{E}_{\pi(\alpha(y^{t-1}))[y_t]}\left[\min_{\mu_i \in D_i(y^{t-1})(y_t)} \mathbb{E}_{\mu_i[\hat{y}]}\left[\tilde{w}_i(y^{t-1}, \hat{y})\right]\right] \\
&= \mathbb{E}_{\pi(\alpha(y^{t-1}))[y_t]}\left[\min_{\mu_i \in D_i(y^{t-1})(y_t)} \mathbb{E}_{\mu_i[\hat{y}]}\left[\min_{\mu_{-i} \in D_{-i}(y^{t-1})(\hat{y})} \mathbb{E}_{\mu_{-i}[\hat{y}]}\left[w_i(y^{t-1}, \hat{y})\right]\right]\right] \\
&= \min_{\{\mu_i \in D_1(y^{t-1})(y_t)\}_{y_t \in Y}} \min_{\{\mu_{-i} \in D_{-i}(y^{t-1})(\hat{y})\}_{\hat{y} \in Y}} \mathbb{E}_{\pi(\alpha(y^{t-1}))[y_t]}\left[\mathbb{E}_{\mu_i[\hat{y}]}\left[\mathbb{E}_{\mu_{-i}[\hat{y}]}\left[w_i(y^{t-1}, \hat{y})\right]\right]\right] \\
&= \min_{\nu \in \mathcal{S}_i'(y^{t-1})} \mathbb{E}_{\nu[y]}[w_i(y^{t-1}, y)].
\end{aligned}
$$

Hence, w from the RDE satisfies equation (O.15). Expanding out the difference equation gives

$$
\begin{aligned}
v_i(y^t) = &\min_{\nu \in \mathcal{S}_i'(y^t)} (1-\beta) g_i(\alpha(y^t)) + \beta \mathbb{E}_{\nu[y]}[w_i(y^t, y)] \\
&\vdots \\
&= \min_{\nu \in \mathcal{S}_i(y^t)} (1-\beta) g_i(\alpha(y^t)) + \mathbb{E}_\nu\left[(1-\beta)\sum_{s=1}^{T-1} \beta^s g_i(\alpha(y^{t+s})) + \beta^T w_i(y^{t+(T-1)}, y^{t+T})\right].
\end{aligned}
$$

Since $w_i(y^{t+s-1}, y)$ is bounded, letting $T \to \infty$ clearly shows that

$$
v_i(y^t) = w_i(y^{t-1}, y_t) = \min_{\nu \in \mathcal{S}_i(y^t)} \left\{(1-\beta) g_i(\alpha(y^t)) + (1-\beta)\mathbb{E}_\nu\left[\sum_{s=1}^{\infty} \beta^s g_i(\alpha(y^{t+s}))\right]\right\}.
$$

Hence the continuation values from the RDE match up with the computed continuation values from the associated PPED. Incentive compatibility and consistency of w as a PPED follow immediately from conditions (i)–(iii) in the definition of RDE. Hence we have shown that an RDE is also a PPED.

Now we show the other direction. Let $v(y^t)$ be the PPED payoffs from (O.13) and let $w(y^{t-1}, y_t)$ be the associated continuation payoffs. All that is required is to show that w satisfies the difference equation in (O.15). Noting that the minimization in (O.13) is recursive, it is clear that

$$
v_i(y^t) = w_i(y^{t-1}, y_t) = \min_{\nu \in \mathcal{S}_i(y^t)} \left\{(1-\beta) \cdot g_i(\alpha(y^t)) + \beta \cdot \mathbb{E}_{\nu[y]}[w_i(y^t, y)]\right\}.[5]
$$

Now by reversing the argument that led us to (O.15), we see that

$$
\begin{aligned}
v_i(y^t) = w_i(y^{t-1}, y_t) &= \min_{\nu \in \mathcal{S}_i(y^t)} \left\{(1-\beta) \cdot g_i(\alpha(y^t)) + \beta \cdot \mathbb{E}_\nu[w_i(y^t, y)]\right\} \\
&= (1-\beta) \cdot g_i(\alpha(y^t)) + \beta \cdot \mathbb{E}_{\pi(\alpha(y^t))}[\tilde{w}_i(y^t, y)].
\end{aligned}
$$

Consistency and incentive compatibility in the definition of RDE follow immediately from the

[5]Note that an element of $\mathcal{S}_i$ gives a distribution over all future public signals, but for the purposes of this equation, we only worry about the distribution it induces over the next period's signal.

corresponding conditions in the definition of PPED. Thus we have shown that a PPED is also an RDE. It follows that the RDE and PPED formulations are equivalent. □

This equivalence helps us develop the APS-style characterization in the next section.

C.2. APS-style Characterization of PPED Payoffs

As in APS, we view a PPED as a pair (α, w) that specifies an action profile α in period 0 along with continuation payoffs $w(y)$ for all signals $y \in Y$. The enforceability condition from APS is modified so that incentives are given by an appropriate within-period minimization.

Definition O.3 (D-enforceable)**.** *The pair (α, v) of an action profile and a payoff is D-enforceable with respect to β and a set W of allowed continuation payoffs if there exists $D \in \mathcal{D}^2$ and a $w(y) \in W$ for each $y \in Y$ such that*

(i) $v_i = (1-\beta)g_i(\alpha) + \beta\mathbb{E}_{\pi(\alpha)[y]}\tilde{w}_i(y)$,

(ii) $\alpha_i \in \arg\max_{\alpha_i'}(1-\beta)g_i(\alpha_i', \alpha_{-i}) + \beta\mathbb{E}_{\pi(\alpha_i', \alpha_{-i})[y]}\tilde{w}_i(y)$, *and*

(iii) the triple $(w, \tilde{w}, D)$ is consistent.

Definition O.4. *For any set $W \subseteq \mathbb{R}^2$,*

$$B_D(\beta, W) \equiv \left\{(1-\beta)g(\alpha) + \beta\mathbb{E}_{\pi(\alpha)[y]}\tilde{w}(y) : (\alpha, v) \text{ is D-enforceable with respect to } W\right\}.$$

We say a set W is *D-self generating (D-SG)* if $W \subseteq B_D(\beta, W)$.

Lemma O.1. *If W is bounded and D-SG, then $W \subseteq E_D(\beta)$.*

Proof. We show that any $v \in W$ can be supported as an RDE. Since $v \in W$, there exists (α, w, D) that D-enforces v. In particular,

$$v = (1-\beta)g(\alpha) + \beta\mathbb{E}_{\pi(\alpha)}\left[\tilde{w}(y)\right]$$

with $w(y) \in W$. Since W is self-generating, for each $v(y_0) = w(y_0)$ there exists $(\alpha(y^0), w(y^0), D(y^0))$ with $w(y_0, y_1) \in Y$ which enforces $v(y_0)$. Iterating this process forward, we will have a sequence $\{(\alpha(y^t), w(y^t), D(y^t))\}_{t\geq 1}$. Since $v(y^{t-1})$ is D-enforced by $(\alpha(y^{t-1}), w(y^{t-1}), D(y^{t-1}))$ and W is a bounded set (so that the $w(y^t)$ are bounded), this exactly shows that the constructed sequence $\{(\alpha(y^t), w(y^t), D(y^t))\}_{t\geq 1}$ is an RDE and thus a PPED by Theorem O.1. □

As in APS, we see that $E_D(\beta)$ is the largest fixed point of the $B_D(\beta, \cdot)$ operator.

Lemma O.2. $E_D(\beta) = B_D(\beta, E_D(\beta))$.

Proof. The fact that $E_D(\beta) \subseteq B_D(\beta, E_D(\beta))$ follows trivially from the recursive structure of the PPED. In particular, after any y_0 the continuation strategies are again a PPED. Thus $v(y_0) \in E_D(\beta)$ for all $y_0 \in Y$. Thus it is clear that $E_D(\beta) \subseteq B_D(\beta, E_D(\beta))$.

Now note that if $v \in B_D(\beta, E_D(\beta))$, then it is enforceable with continuation payoffs that are from PPEDs and thus is clearly a PPED. Thus, $E_D(\beta) \supseteq B_D(\beta, E_D(\beta))$ so that we have the equivalence. □

C.3. Proof of Theorem 4 (Characterizing the Limit PPED Set)

Proof. The proof of this theorem follows that of Theorem 3.1 in FL very closely. First, we prove that $E_D(\beta) \subseteq Q_P$. Suppose it does not hold. Then, we can find $\lambda \in \mathbb{R}^2$, $\beta \in (0,1)$, and $v \in E_D(\beta)$ so that $\lambda \cdot v \equiv k > k^*_A(\lambda)$. Let $E^*_D(\beta)$ be the convex hull of $E_D(\beta)$. Let $\overline{k}$ be the maximal score in $E^*_D(\beta)$ in direction λ. Note that $\overline{k} > k^*_P(\beta)$. Since the score function is linear, there exists $v \in E_D(\beta)$ with $\lambda \cdot v = \overline{k}$. Then by the fact that $E_D(\beta) \subseteq B_D(\beta, E_D(\beta))$, we know that v is enforceable with continuation values $w(y)$ in $E_D(\beta)$. But, the score $\overline{k}$ is maximal, so we know that $\lambda \cdot v \geq \lambda \cdot w(y)$. Hence the value of the programming problem is at least $\overline{k} > k^*_P(\beta)$, which is a contradiction.

We now move to the second claim. Suppose Q_P has full dimension and pick a smooth convex set W in the interior of Q_P. It is sufficient to prove that for any such set W and any $v \in W$ there is an open set $U \ni v$ and a β such that $U \subseteq B_D(\beta, W)$. This is easy if point v lies in the interior of W (by using a static Nash equilibrium), so focus on v on the boundary of W. Let λ be the unique vector orthogonal to the tangent of W at v, let $H(\lambda, k)$ be the half-space $\{v' : \lambda \cdot v' \leq k\}$, and let α be an action profile which enforces the score $k^*(\lambda)$ in direction λ. Then, for some β' and $\epsilon > 0$, (α, v) can be enforced with respect to $H(\lambda, k - \epsilon)$; that is, in particular,

$$v_i = (1 - \beta')g_i(\alpha) + \beta' \mathbb{E}_{\pi(\alpha)[y]} \tilde{w}_i(y),$$

with $w(y) \in H(\lambda, k - \epsilon)$. This holds since the map from w to $\tilde{w}$ is translation-preserving.[6] Thus, for all $\beta'' \geq \beta'$, (α, v) can be enforced with respect to the half-plane $H(\lambda, k - \beta'(1 - \beta'')/\beta''(1 - \beta')\epsilon)$ so that all continuation values are in an $\kappa(1 - \beta'')$-ball around v (for some constant $\kappa > 0$), because the map from w to $\tilde{w}$ is scaling-invariant. As W is smooth, this means that there is some large $\beta < 1$ so that (α, v) is enforceable with respect to continuation values in the interior of W. Moving around the continuation values in a small neighborhood yields the desired open neighborhood U containing v. □

C.4. Proof of Theorem 5 (Anti-Folk Theorem)

Proof. Suppose there existed an action profile a with $g(a) \in Q_D$ and $g(a)$ on the Pareto frontier of V^*. Then, a has full score in some direction $\lambda \gg 0$, that is, $k_D(a, \lambda) = \lambda \cdot g(a)$. By the definition of k_D in (4), this requires that all continuation values $\tilde{w}(y)$ lie on the hyperplane defined by λ, so $\lambda \cdot \tilde{w}(y) = \lambda \cdot g(a)$ for all y. This, however, is impossible for values $\tilde{w}(y)$ that lie on a positively sloped line, unless the values are equal, i.e. $\tilde{w}(y) = \tilde{w}(y')$ for all y and y'. If continuation values were equal, then they could not provide any incentives to enforce a, so a would have to be a stage game Nash equilibrium. This was ruled out in the statement of Theorem 5 and completes the proof. □

C.5. Proof of Theorem 6

We first formally define k_L and Q_L.

[6] By "translation-preserving" we mean that if $(w, \tilde{w}, D)$ is a consistent triple, then $(w + e, \tilde{w} + e, D)$ is also a consistent triple, where $w + e$ denotes adding the ordered pair (e_1, e_2) to each $w(y)$. The only thing one needs to do here is shift the set of optimal w's from the score problem $k^*(\lambda)$ a little bit towards W until v is enforceable. The $\tilde{w}$'s merely "follow" the w's.

Definition O.5 (Linearly Enforceable). *A pair (α, v) is said to be linearly enforceable with respect to (β, W) if there exists $w(y) \in W$ lying on a positively sloped line such that*

(i) $v = (1-\beta)g(\alpha) + \beta\mathbb{E}_{\pi(\alpha)[y]}[w(y)]$ *and*

(ii) $\alpha_i \in \arg\max_{\alpha'_i} \left\{(1-\beta)g(\alpha'_i, \alpha_{-i}[y]) + \beta\mathbb{E}_{\pi(\alpha'_i,\alpha_{-i})}[w(y)]\right\}$.

For any set $W \subseteq \mathbb{R}^2$, let

$$B_L(\beta, W) \equiv \left\{v = (1-\beta)g(\alpha) + \beta\mathbb{E}_{\pi(\alpha)[y]}w(y) : (\alpha, v) \text{ is linearly enforceable with respect to } W\right\}.$$

A set $W \subseteq \mathbb{R}^2$ is said to be *linearly self-generating* if $W \subseteq B_L(\beta, W)$. Let $E_L(\beta)$ be the set of payoffs of totally linear PPEs. Then, analogously to Lemma O.1, we have

Lemma O.3. *If W is bounded and linearly self-generating then $W \subseteq E_L(\beta)$.*

As before, we also have $E_L(\beta) = B_L(\beta, E_L(\beta))$ so that $E_L(\beta)$ is the largest fixed point of the B_L operator. Using this APS characterization, we get a corresponding programming problem. For a strategy profile α and a direction $\lambda \in \mathbb{R}^2$, define $k_L(\alpha, \lambda, \beta)$ as the value of the program

$$\begin{aligned} &\sup_{v,w(y)} \; \lambda \cdot v \\ \text{s.t.} \quad & v = (1-\beta)g(\alpha) + \beta\mathbb{E}_{\pi(\alpha)[y]}w(y) \\ & v_i = (1-\beta)g_i(a_i, \alpha_{-i}) + \beta\mathbb{E}_{\pi(a_i,\alpha_{-i})[y]}w_i(y) \quad \forall a_i \in \operatorname{supp}\alpha_i \\ & v_i \geq (1-\beta)g_i(a_i, \alpha_{-i}) + \beta\mathbb{E}_{\pi(a_i,\alpha_{-i})[y]}w_i(y) \quad \forall a_i \in A_i \\ & \lambda \cdot v \geq \lambda \cdot w(y) \qquad \forall y \in Y \\ & w(y) \text{ lies on a line with slope in } (0,\infty). \end{aligned} \tag{O.16}$$

Define $k_L^*(\lambda, \beta) \equiv \sup_\alpha k_L(\alpha, \lambda, \beta)$.[7] As before $k_L(\alpha, \lambda, \beta)$ is independent of β by a standard scaling argument. We can show as before that if Q_L is defined analogously to Q_D in Theorem 4, i.e.,

$$Q_L \equiv \bigcap_{\lambda \in \mathbb{R}^2} \left\{v \in \mathbb{R}^2 : \lambda \cdot v \leq k_L^*(\lambda)\right\}, \tag{O.17}$$

then the analogous result to Theorem 4 holds as well.

We now prove Theorem 6. We do this in three steps. First, we prove that $E_L(\beta) \subseteq E_D(\beta)$, then that $Q_L \subseteq Q_D$, and finally that for any payoff profile $v_{NE} \in \mathbb{R}^2$ of the stage game, $Q_D \cap \{v \geq v_{NE}\} \subseteq Q_L$.

[7] Programs (4) and (O.16) may bear resemblance to the program suggested in Fudenberg, Levine, and Takahashi (2007) (FLT), since we have an additional constraint on the continuation payoffs. Note, however, that the results in this section are not simply applications of FLT. FLT computes the limit set of payoffs from $\mathcal{A}^0$-PPEs, which place constraints on the actions allowed in equilibrium after any public history. The programs in our setup place restrictions on the *continuation payoffs* that do not map to simple restrictions on actions. A difference between totally linear PPEs and strongly symmetric PPEs is important here. The method in FLT applies to strongly symmetric PPEs since the same restriction on actions can be placed after each public history; for totally linear PPEs, since the slope and intercept of the continuation values can be different at each history, there is no such restriction on actions that applies globally after each public history.

Step 1. Choose $v \in E_L(\beta)$ and let σ be a linear PPE with payoff v. Let α be the action in the first period. Let $w(y)$ be the continuation payoffs after signal y realizes. Then note that $w(y) \in E_L(\beta)$ since any subgame of a linear PPE is also a linear PPE. Now choose $\tilde{w}(y) = w(y)$ and back out the correspondence D from the consistency requirement so that $(w, \tilde{w}, D)$ are consistent. Finally, since no player has an incentive to deviate in the first period of σ we know that

$$\alpha_i \in \arg\max_{a_i} (1-\beta) g_i(a_i, \alpha_{-i}) + \beta E_{\pi(a_i,\alpha_{-i})[y]}[w_i(y)].$$

Now, since $w \equiv \tilde{w}$, this is exactly says that (α, v) is D-enforceable by $(w, \tilde{w}, D)$. This shows that $E_L(\beta) \subseteq E_D(\beta)$.

Step 2. To show $Q_L \subseteq Q_P$, we simply compare Programs (4) and (O.16). Fix α. Let $(v, \{w(y)\})$ be feasible in (O.16). Then, setting $\tilde{w}(y) = w(y)$ and determining D from (3) gives us a feasible set $(v, \{w(y)\}, D)$ in (4) since $w(y)$ already lie on a line of positive slope. Thus, $k_D(\alpha, \lambda, \beta) \geq k_L(\alpha, \lambda)$ and thus $k_D^*(\lambda) \geq k_L^*(\lambda)$ for all λ, meaning $Q_L \subseteq Q_D$.

Step 3. We first prove the following helpful lemma.

Lemma O.4. *If $(w, \tilde{w}, D)$ is consistent, $w(y) \in W$ for all $y \in Y$, and W is convex and meet-closed (i.e., the meet of any set of points in W is also in W), then $\tilde{w}(y) \in W$ for all $y \in Y$.*

Proof. We already know that $\{\tilde{w}(y)\}$ is linear with a positive slope. Hence the points are Pareto-ranked. If $\{\tilde{w}(y)\}$ is a singleton then $D_i(y) = \Delta(Y)$ for all $y \in Y$, and thus $\tilde{w}(y) = \bigwedge\{w(y)\}_{y \in Y} \in W$. Now suppose $\{\tilde{w}(y)\}$ has at least two distinct points. Let $\tilde{w}(y_0)$ be the Pareto-worst point. Then we know that $D_i(y_0) = \Delta(Y)$, so that we still have

$$\tilde{w}(y_0) = \bigwedge\{w(y)\}_{y \in Y} \in W,$$

since W was assumed to be meet-closed. Now let $\tilde{w}(y_1)$ be the Pareto-best point, and let $Y' \subset Y$ be the set of signals such $\tilde{w}(y) = \tilde{w}(y_1)$ for only $y \in Y'$. Then we know that $D_i(y) = \Delta(Y')$ for all $y \in Y'$. But then using the definition of $\tilde{w}$ we exactly have

$$\tilde{w}(y') = \bigwedge\{w(y)\}_{y \in Y'} \in W$$

when $y' \in Y$. Thus, $\tilde{w}(y_1) \in W$ as well. All other points of $\tilde{w}$ lie in a line between $\tilde{w}(y_0)$ and $\tilde{w}(y_1)$, both of which lie in W. Since W is convex, we have the result. □

Now we prove Step 3. Suppose either $\lambda_1 \geq 0$ or $\lambda_2 \geq 0$. Fix α and consider a triple $(v, \{w(y)\}, D)$ that is feasible in Program (4). Since the half plane $H(\lambda, \lambda \cdot v) \equiv \{w' \in \mathbb{R}^2 : \lambda \cdot w' \leq \lambda \cdot v\}$ is meet-closed, Lemma O.4 says that $\tilde{w}(y) \in H(\lambda, \lambda \cdot v)$ for all y as well. It follows that $(v, \{\tilde{w}(y)\})$ is feasible in Program (O.16), meaning $k_L(\alpha, \lambda) \geq k_D(\alpha, \lambda)$, or $k_L^*(\lambda) \geq k_D^*(\lambda)$. For $\lambda_1 < 0$ and $\lambda_2 < 0$, we know that $k_L^*(\lambda) \geq v_{NE}$, since v_{NE} is supported by a Nash equilibrium.

To conclude the proof of Theorem 6, note that

$$\begin{aligned}\mathcal{L}(v_{NE}) &\subseteq \bigcap_{\{\lambda:\lambda_1<0 \text{ and } \lambda_2<0\}} H(\lambda, \lambda\cdot v_{NE}) \cap \bigcap_{\{\lambda:\lambda_1\geq 0 \text{ or } \lambda_2\geq 0\}} H(\lambda, k_D^*(\lambda)) \\ &\subseteq \bigcap_{\lambda} \left\{w' \in \mathbb{R}^2 : \lambda\cdot w' \leq k_D^*(\lambda)\right\} = Q_L.\end{aligned}$$

References

Fudenberg, D., D. K. Levine, and S. Takahashi (2007): "Public Perfect Equilibrium When Players Are Patient," *Games and Economic Behavior*, 61(1), 27–49.

Kreps, D. M., and R. Wilson (1982): "Sequential Equilibria," *Econometrica*, 50(4), 863–894.