Econometrica Supplementary Material

SUPPLEMENT TO "PROGRESSIVE LEARNING": ADDITIONAL APPENDIXES

(*Econometrica*, Vol. 85, No. 6, November 2017, 1965–1990)

AVIDIT ACHARYA
Department of Political Science, Stanford University

JUAN ORTNER
Department of Economics, Boston University

These Appendixes present the proof of Lemma 0, an extension of our equilibrium characterization allowing for mixed strategies, an analysis of the full commitment case, and details for Example 4 showing that path dependence can arise when shocks are ergodic.

KEYWORDS: Principal–agent model, adverse selection, ratchet effect, inefficiency, learning, path dependence.

APPENDIX SA.1: PROOF OF LEMMA 0

PROOF OF PART (I): The proof is by strong induction on the cardinality of the support of the principal's beliefs, $C[h_t]$. Fix an equilibrium (σ, μ) and note that the claim is true for all histories h_t such that $|C[h_t]| = 1$.[1] Suppose next that the claim is true for all histories h with $|C[h]| \le n-1$ and consider a history h_t with $|C[h_t]| = n$.

Suppose by contradiction that $V^{(\sigma,\mu)}_{\overline{k}[h_t]}[h_t, b_t] > 0$. Then there must exist a history $(h_{t'}, b_{t'})$ with $h_{t'} \succeq h_t$ that arises on the path of play with positive probability at which the principal offers a transfer $T_{t'} > c_{\overline{k}[h_t]}$ that type $c_{\overline{k}[h_t]}$ accepts. Note first that since type $c_{\overline{k}}[h_t]$ accepts offer $T_{t'}$, all types in the support of $C[h_{t'}]$ must also accept it. Indeed, if this were not true, then there would be a highest type $c_k \in C[h_{t'}]$ that rejects the offer. By the induction hypothesis, the equilibrium payoff that this type obtains at history $h_{t'}$ is $V^{(\sigma,\mu)}_k[h_{t'}, b_{t'}] = 0$, since this type would be the highest cost in the support of the principal's beliefs following a rejection. But this cannot be, since type c_k can get a payoff of at least $T_{t'} - c_k > 0$ by accepting the principal's offer at time t'.

We now construct an alternative strategy profile $\tilde{\sigma}$ that is otherwise identical to σ except that at history $(h_{t'}, b_{t'})$, the agent is offered a transfer $\tilde{T} \in (c_{\overline{k}[h_t]}, T_{t'})$. Specify the principal's beliefs at history $(h_{t'}, b_{t'})$ as follows: regardless of the agent's action, the principal's beliefs at the end of the period are the same as her beliefs at the beginning of the period. At all other histories, the principal's actions and beliefs are the same as in the original equilibrium. Note that, given these beliefs, at history $h_{t'}$, all agent types in $C[h_{t'}]$ find it strictly optimal to accept the principal's offer $\tilde{T}$ and take the action. Thus, the principal's payoff at history $h_{t'}$ is larger than her payoff under the original equilibrium, which cannot be since the original equilibrium was in Σ_K. *Q.E.D.*

PROOF OF PART (II): The proof is by induction of the cardinality of $C[h_t]$. Consider first a history h_t such that $|C[h_t]| = 1$. Since $|C[h_t]| = 1 < 2$, the claim is vacuously true.

Avidit Acharya: avidit@stanford.edu
Juan Ortner: jortner@bu.edu

[1] Indeed, if $C[h_t] = \{c_i\}$, then in any PBE in Σ_K the agent takes action $a = 1$ at time $t' \ge t$ if and only if $b_{t'} \in E_i$, and the principal pays the agent a transfer equal to c_i every time the agent takes the action.

 DOI: 10.3982/ECTA14718

Suppose next that the result holds for all histories h such that $|C[h]| \leq n-1$ and consider a history h_t such that $|C[h_t]| = n$. Consider two "adjacent" types $c_i, c_{i+1} \in C[h_t]$. We have two possible cases: (i) with probability 1, types c_i and c_{i+1} take the same action at all histories $(h_{t'}, b_{t'})$ with $h_{t'} \succeq h_t$; (ii) there exists a history $(h_{t'}, b_{t'})$ with $h_{t'} \succeq h_t$ at which types c_i and c_{i+1} take different actions. Under case (i),

$$\begin{aligned} V_i^{(\sigma,\mu)}[h_t, b_t] &= \mathbb{E}^{(\sigma,\mu)}\left[\sum_{t'=t}^{\infty} \delta^{t'-t}(1-\delta)(T_{t'} - c_i)a_{t',i}\middle| h_t, b_t\right] \\ &= \mathbb{E}^{(\sigma,\mu)}\left[\sum_{t'=t}^{\infty} \delta^{t'-t}(1-\delta)(T_{t'} - c_{i+1})a_{t',i+1}\middle| h_t, b_t\right] \\ &\quad + \mathbb{E}^{(\sigma,\mu)}\left[\sum_{t'=t}^{\infty} \delta^{t'-t}(1-\delta)(c_{i+1} - c_i)a_{t',i+1}\middle| h_t, b_t\right] \\ &= V_{i+1}^{(\sigma,\mu)}[h_t, b_t] + A_{i+1}^{(\sigma,\mu)}[h_t, b_t](c_{i+1} - c_i). \end{aligned}$$

For case (ii), let $\underline{t} = \min\{t' \geq t : a_{t',i+1} \neq a_{t',i}\}$ be the first time after t at which types c_i and c_{i+1} take different actions. Let $c_k \in C[h_{\underline{t}}]$ be the highest-cost type that takes the action at time $\underline{t}$. The transfer $T_{\underline{t}}$ that the principal offers at time $\underline{t}$ must satisfy $V_k^{(\sigma,\mu)}[h_{\underline{t}}, b_{\underline{t}}] = (1-\delta)(T_{\underline{t}} - c_k) = V_{k+1}^{(\sigma,\mu)}[h_{\underline{t}}, b_{\underline{t}}] + A_{k+1}^{(\sigma,\mu)}[h_{\underline{t}}, b_{\underline{t}}](c_{k+1} - c_k)$.[2] Note further that $V_{k+1}^{(\sigma,\mu)}[h_{\underline{t}}, b_{\underline{t}}] \geq (1-\delta)(T_{\underline{t}} - c_{k+1})$, since an agent with cost c_{k+1} can guarantee $(1-\delta)(T_{\underline{t}} - c_{k+1})$ by taking the action at time $\underline{t}$ and then not taking the action in all future periods. Since $(1-\delta)(T_{\underline{t}} - c_k) = V_{k+1}^{(\sigma,\mu)}[h_{\underline{t}}, b_{\underline{t}}] + A_{k+1}^{(\sigma,\mu)}[h_{\underline{t}}, b_{\underline{t}}](c_{k+1} - c_k)$, it follows that $A_{k+1}^{(\sigma,\mu)}[h_{\underline{t}}, b_{\underline{t}}] \leq 1-\delta$.

We now show that all types below c_k also take the action at time $\underline{t}$. That is, we show that all agents in the support of $C[h_{\underline{t}}]$ with cost weakly lower than c_k take the action at $\underline{t}$, and all agents with cost weakly greater than c_{k+1} do not take the action. Note that this implies that $c_i = c_k$ (since types c_i and c_{i+1} take different actions at time $\underline{t}$). Suppose for the sake of contradiction that this is not true and let c_j be the highest-cost type below c_k that does not take the action. The payoff that this agent gets from not taking the action is $V_{j\to k+1}^{(\sigma,\mu)}[h_{\underline{t}}, b_{\underline{t}}] = V_{k+1}^{(\sigma,\mu)}[h_{\underline{t}}, b_{\underline{t}}] + A_{k+1}^{(\sigma,\mu)}[h_{\underline{t}}, b_{\underline{t}}](c_{k+1} - c_j)$, which follows since at time $\underline{t}$, types c_j and c_{k+1} do not take the action and since, by the induction hypothesis, from time $\underline{t}+1$ onward, the payoff that an agent with cost c_j gets is equal to what this agent would get by mimicking an agent with cost c_{k+1}. On the other hand, the payoff that agent c_j obtains by taking the action and mimicking type c_k is

$$\begin{aligned} V_{j\to k}^{(\sigma,\mu)}[h_{\underline{t}}, b_{\underline{t}}] &= V_k^{(\sigma,\mu)}[h_{\underline{t}}, b_{\underline{t}}] + A_k^{(\sigma,\mu)}[h_{\underline{t}}, b_{\underline{t}}](c_k - c_j) \\ &= (1-\delta)(T_{\underline{t}} - c_j) + A_k^{(\sigma,\mu)}[h_{\underline{t}}, b_{\underline{t}}](c_k - c_j) \\ &= V_{k+1}^{(\sigma,\mu)}[h_{\underline{t}}, b_{\underline{t}}] + A_{k+1}^{(\sigma,\mu)}[h_{\underline{t}}, b_{\underline{t}}](c_{k+1} - c_k) + A_k^{(\sigma,\mu)}[h_{\underline{t}}, b_{\underline{t}}](c_k - c_j) \\ &> V_{k+1}^{(\sigma,\mu)}[h_{\underline{t}}, b_{\underline{t}}] + A_{k+1}^{(\sigma,\mu)}[h_{\underline{t}}, b_{\underline{t}}](c_{k+1} - c_j), \end{aligned}$$

[2]The first equality follows since, after time $\underline{t}$, type c_k is the highest type in the support of the principal's beliefs if the agent takes action $a = 1$ at time $\underline{t}$. The second equality follows since we focus on PBE in Σ_K, so the transfer $T_{\underline{t}}$ leaves a c_k agent indifferent between accepting and rejecting.

where the inequality follows since $A_{k+1}^{(\sigma,\mu)}[h_{\underline{t}}, b_{\underline{t}}] \leq 1-\delta < A_k^{(\sigma,\mu)}[h_{\underline{t}}, b_{\underline{t}}]$.[3] Hence, type j strictly prefers to take the action, a contradiction. Therefore, all types below c_k take the action at time $\underline{t}$ and so $c_i = c_k$.

By the arguments above, the payoff that type $c_i = c_k$ obtains at time $\underline{t}$ is

$$V_i^{(\sigma,\mu)}[h_{\underline{t}}, b_{\underline{t}}] = (1-\delta)(T_{\underline{t}} - c_i) = V_{i+1}^{(\sigma,\mu)}[h_{\underline{t}}, b_{\underline{t}}] + A_{i+1}^{(\sigma,\mu)}[h_{\underline{t}}, b_{\underline{t}}](c_{i+1} - c_i),$$

since the transfer that the principal offers at time $\underline{t}$ satisfies $(1-\delta)(T_{\underline{t}} - c_i) = V_{i+1}^{(\sigma,\mu)}[h_{\underline{t}}, b_{\underline{t}}] + A_{i+1}^{(\sigma,\mu)}[h_{\underline{t}}, b_{\underline{t}}](c_{i+1} - c_i)$. Moreover,

$$\begin{aligned} V_i^{(\sigma,\mu)}[h_t, b_t] &= \mathbb{E}^{(\sigma,\mu)}\left[\sum_{t'=t}^{\underline{t}-1} \delta^{t'-t}(1-\delta)(T_{t'} - c_i)a_{t',i} + \delta^{\underline{t}-t} V_i^{(\sigma,\mu)}[h_{\underline{t}}, b_{\underline{t}}] \middle| h_t, b_t\right] \\ &= \mathbb{E}^{(\sigma,\mu)}\left[\sum_{t'=t}^{\underline{t}-1} \delta^{t'-t}\big((1-\delta)(T_{t'} - c_{i+1})a_{t',i+1} + (1-\delta)(c_{i+1} - c_i)a_{t',i+1}\big) \middle| h_t, b_t\right] \\ &\quad + \mathbb{E}^{(\sigma,\mu)}\big[\delta^{\underline{t}-t}\big(V_{i+1}^{(\sigma,\mu)}[h_{\underline{t}}, b_{\underline{t}}] + A_{i+1}^{(\sigma,\mu)}[h_{\underline{t}}, b_{\underline{t}}](c_{i+1} - c_i)\big) | h_t, b_t\big] \\ &= V_{i+1}^{(\sigma,\mu)}[h_t, b_t] + A_{i+1}^{(\sigma,\mu)}[h_t, b_t](c_{i+1} - c_i), \end{aligned}$$

where the second equality follows since $a_{t',i} = a_{t',i+1}$ for all $t' \in \{t, \dots, \underline{t}-1\}$. Hence, the result also holds for histories h_t with $|C[h_t]| = n$. *Q.E.D.*

APPENDIX SA.2: MIXED STRATEGIES

This appendix extends the results in the main text to allow for mixed strategies. In particular, we show that the equilibrium we characterize in Theorem 1 remains the unique PBE that is sequentially optimal for the principal among all finitely revealing PBE; that is, among all PBE in which, along any history, the principal updates her beliefs a finite number of periods.

Fix a PBE (σ, μ), with $\sigma = (\tau, \{\alpha_k\}_{k=1}^K)$. For any history (h_t, b_t), we say that period t is a period of revelation if (a) $\mu[h_t] \notin S_1$ (i.e., if the principal is uncertain about the agent's type) and (b) there exists $c_i, c_j \in C[h_t]$ such that $\alpha_i(h_t, b_t) \neq \alpha_j(h_t, b_t)$ (i.e., there exist at least two types in the support of the principal's beliefs that play different—possibly mixed—actions at history (h_t, b_t)). We say that an equilibrium (σ, μ) is T-revealing if, for any t and along any history h_t, the number of periods of revelation $t' < t$ is not greater than T.[4]

Three things are worth noting about T-revealing PBE. First, a T-revealing strategy does not put any bound on the occurrence of the last period of revelation. Hence, information may be revealed at any point during the game. Second, a T-revealing strategy does not require the agent to reveal her information fully. Third, since the set of possible types of the agent is finite, any pure strategy PBE is T-revealing for some T.

[3]Recall that, for all (h_t, b_t), $A_k^{(\sigma,\mu)}[h_t, b_t] = (1-\delta)\mathbb{E}^{(\mu,\sigma)}[\sum_{t'=t}^{\infty} \delta^{t'-t} a_{t',k} | b_t, h_t]$. By assumption, an agent with type c_k takes action $a = 1$ at time $\underline{t}$, so $a_{\underline{t},k} = 1$. Moreover, it is easy to show that an agent with cost c_k will take action $a = 1$ with positive probability at some date $t > \underline{t}$. Therefore, $A_k^{(\sigma,\mu)}[h_{\underline{t}}, b_{\underline{t}}] > 1-\delta$.

[4]This definition is borrowed from Peski (2008).

Let Σ_0^M denote the set of PBE that are finitely revealing (i.e., the set of PBE that are T-revealing for some finite T). For all $k = 1, \dots, K$, we define the sets Σ_k^M recursively as

$$\Sigma_k^M := \left\{ (\sigma, \mu) \in \Sigma_{k-1}^M : \begin{array}{l} \forall (h_t, b_t) \text{ with } \mu[h_t] \in S_k \text{ and } \forall (\sigma', \mu') \in \Sigma_{k-1}^M, \\ U^{(\sigma,\mu)}[h_t, b_t] \geq W^{(\sigma',\mu')}[\mu[h_t], b_t] \end{array} \right\}.$$

Let (σ^P, μ^P) denote the PBE characterized in Theorem 1 and note that $(\sigma^P, \mu^P) \in \Sigma_0^M$. The following theorem shows that (σ^P, μ^P) belongs to the set Σ_K^M. Note that this implies that any PBE in Σ_K^M gives the principal the same payoff as (σ^P, μ^P) at every history. Moreover, as the proof the theorem clarifies, any equilibrium $(\sigma, \mu) \in \Sigma_K^M$ induces the same outcome as (σ^P, μ^P).

THEOREM SA.1: *We have* $(\sigma^P, \mu^P) \in \Sigma_K^M$.

PROOF: Fix a finitely revealing equilibrium $(\sigma, \mu) \in \Sigma_K^M$ and let T be the upper bound on the periods of revelation under (σ, μ). We start by showing that, at histories at which there have already been T periods of information revelation, players' behavior under (σ, μ) must coincide with their behavior under (σ^P, μ^P).

Consider a history (h_t, b_t) at which there have already been T periods of information revelation. Hence, $\mu[h_t] = \mu[h_{t+s}]$ for all $s \geq 0$ and all histories h_{t+s} that follow history h_t. This implies that

$$U^{(\sigma,\mu)}[h_t, b_t] \leq (1-\delta)\mathbb{E}\left[\sum_{s=0}^{\infty} \delta^s \mathbf{1}_{\{b_{t+s} \in E_{\overline{k}[h_t]}\}} (b_{t+s} - c_{\overline{k}[h_t]}) \middle| b_t \right], \tag{SA.1}$$

where $U^{(\sigma,\mu)}[h_t, b_t]$ is the principal's continuation payoff at history (h_t, b_t). To see why the inequality holds, note that all agent types in the support of $\mu[h_t]$ use the same strategy at all periods after time t. Moreover, since an agent of type $c_{\overline{k}[h_t]}$ gets a continuation payoff of 0 at all histories, she only takes the action at time $\tau \geq t$ if $T_\tau = c_{\overline{k}[h_t]}$.[5] These two observations together imply the bound in equation (SA.1). Since the principal's continuation payoff at history (h_t, b_t) under equilibrium (σ^P, μ^P) is weakly larger than the right-hand side of (SA.1), it follows that players' behavior under (σ, μ) must coincide with their behavior under (σ^P, μ^P) at all histories after information revelation has stopped.

Next, consider a history h_t with the property that, for all histories h_{t+s} with $s \geq 1$ that follow history (h_t, b_t), players' behavior under $(\sigma, \mu) \in \Sigma_K^M$ coincides with their behavior under (σ^P, μ^P). We now show that at such a history (h_t, b_t), the players' behavior under $(\sigma, \mu) \in \Sigma_K^M$ coincides with their behavior under (σ^P, μ^P). Before presenting its proof, we note that this result and the result above together establish Theorem SA.1.

To see why the result is true, we consider two separate cases: (i) b_t such that $X(b_t, E_{\overline{k}[h_t]}) > 1 - \delta$ and (ii) b_t such that $X(b_t, E_{\overline{k}[h_t]}) \leq 1 - \delta$.

Case (i). Let T_t be the principal's offer at history (h_t, b_t) and note that $T_t \leq c_{\overline{k}[h_t]}$ (see footnote 5). We start by showing that if $T_t \leq c_{\overline{k}[h_t]}$ is such that an agent with type $c_{\overline{k}[h_t]}$ rejects the offer with probability 1, then all agent types also reject the offer with probability 1. Suppose by contradiction that the set of types that accept offer T_t with positive probability is nonempty. Let $c_i < c_{\overline{k}[h_t]}$ be the highest cost of a type that accepts

[5] In any PBE in Σ_K^M, the principal never makes an offer T_t that is larger than the highest cost in the support of her beliefs. Indeed, if $T_t > c_{\overline{k}[h_t]}$ for some history (h_t, b_t), we can construct an alternative finitely revealing equilibrium in Σ_{k-1}^M (where $k = |C[h_t]|$) that gives the principal strictly more profits than (σ, μ), which would contradict $(\sigma, \mu) \in \Sigma_k^M$.

T_t with positive probability. The payoff that type c_i obtains by accepting the offer is $(1-\delta)(T_t - c_i) + \delta \times 0 \le (1-\delta)(c_{\overline{k}[h_t]} - c_i)$, since from $t+1$ onward type c_i would be the highest type in the support of the principal's beliefs following an acceptance, and since equilibrium (σ, μ) coincides with (σ^P, μ^P) at all histories that follow history (h_t, b_t). In contrast, the payoff that type c_i gets by rejecting the offer and mimicking type $c_{\overline{k}[h_t]}$ at all times $\tau > t$ is $X(b_t, E_{\overline{k}[h_t]})(c_{\overline{k}[h_t]} - c_i) > (1-\delta)(c_{\overline{k}[h_t]} - c_i)$, a contradiction. Hence, if $T_t \le c_{\overline{k}[h_t]}$ is such that an agent with type $c_{\overline{k}[h_t]}$ rejects the offer with probability 1, then all agent types also reject the offer with probability 1.

There are two subcases to consider: (a) $b_t \in E_{\overline{k}[h_t]}$, and (b) $b_t \notin E_{\overline{k}[h_t]}$. Consider case (a). We show that, in this case, the principal makes offer $T_t = c_{\overline{k}[h_t]}$ and that this offer is accepted by all types with probability 1 (so behavior under equilibrium (σ, μ) coincides with behavior under (σ^P, μ^P)). As a first step, we show that the principal makes offer $T_t = c_{\overline{k}[h_t]}$ and that this offer is accepted by an agent of type $c_{\overline{k}[h_t]}$ with positive probability. Indeed, if this was not the case, then by the arguments above, no agent type would accept offer T_t, so $\mu[h_{t+1}] = \mu[h_t]$. But then we would be able to construct an alternative finitely revealing equilibrium in Σ^M_{k-1} (where $k = |C[h_t]|$) that gives the principal strictly more profits than (σ, μ), which would contradict $(\sigma, \mu) \in \Sigma^M_k$. To see how, consider an equilibrium in which players' behavior is identical to their behavior under (σ, μ) at every history except for history (h_t, b_t). At history (h_t, b_t), the principal makes offer $T_t = c_{\overline{k}[h_t]}$ and every type accepts this offer with probability 1. The principal's beliefs at $t+1$ are identical to $\mu[h_t]$ regardless of whether the agent accepts the offer or not. One can check that this modified strategy profile is a PBE in finitely revealing strategies that lies in Σ^M_{k-1}. Moreover, it delivers the principal a strictly larger payoff at history (h_t, b_t) than (σ, μ), which contradicts $(\sigma, \mu) \in \Sigma^M_K$.

Next we show that offer $T_t = c_{\overline{k}[h_t]}$ is accepted with probability 1 by all agent types $c_i < c_{\overline{k}[h_t]}$. Towards a contradiction, let c_i be the highest-cost type below $c_{\overline{k}[h_t]}$ that rejects the offer. The payoff that this type obtains by rejecting is at most $X(b_t, E_{\overline{k}[h_t]})(c_{\overline{k}[h_t]} - c_i)$, since either type c_i will be the second highest cost in the support of $\mu[h_{t+1}]$ (and type $c_{\overline{k}[h_t]}$ will be the highest) or type c_i will be the highest cost in the support of $\mu[h_{t+1}]$. In contrast, by accepting the offer and then mimicking type $c_{\overline{k}[h_t]}$, she obtains $(1 - \delta + X(b_t, E_{\overline{k}[h_t]}))(c_{\overline{k}[h_t]} - c_i)$, which cannot be. Hence, offer $T_t = c_{\overline{k}[h_t]}$ is accepted with probability 1 by all agent types $c_i < c_{\overline{k}[h_t]}$.

Finally, we show that $T_t = c_{\overline{k}[h_t]}$ is accepted by an agent with cost $c_{\overline{k}[h_t]}$ with probability 1. Suppose by contradiction this is not true, and consider an alternative finitely revealing equilibrium such that players' behavior coincides with their behavior under (σ, μ) at all histories except (h_t, b_t). At such a history, the principal makes offer $T_t = c_{\overline{k}[h_t]}$ and this offer is accepted by all types of the agent with probability 1 (after which the principal's beliefs remain equal to $\mu[h_t]$ regardless of the agent's action). One can check that this is a PBE in Σ_{k-1} and that this PBE gives the principal a strictly larger profit than the original equilibrium (σ, μ), a contradiction. Hence, offer $T_t = c_{\overline{k}[h_t]}$ is accepted by all agent types with probability 1.

Consider next case (b). We show that, in this case, the principal makes an offer $T_t < c_{\overline{k}[h_t]}$ that all agent types reject. From our arguments above, if $T_t \le c_{\overline{k}[h_t]}$ is rejected by an agent of type $c_{\overline{k}[h_t]}$ with probability 1, then the offer is rejected by all agent types $c_i < c_{\overline{k}[h_t]}$ with probability 1. This implies that any offer $T_t < c_{\overline{k}[h_t]}$ is rejected by every agent type with probability 1. Note that in an equilibrium $(\sigma, \mu) \in \Sigma^M_K$, at such a history the principal would never make an offer $T_t = c_{\overline{k}[h_t]}$ that is accepted by an agent of type $c_{\overline{k}[h_t]}$ with positive probability. If this were the case, and by the same arguments used in case (a), such an offer would be accepted by all types $c_i < c_{\overline{k}[h_t]}$ with probability 1. Since $b_t < c_{\overline{k}[h_t]}$, the principal would be strictly better off by making an offer $T_t < c_{\overline{k}[h_t]}$ that is rejected by all types with probability 1.[6]

Case (ii). Consider next histories (h_t, b_t) with $X(b_t, E_{\overline{k}[h_t]}) \leq 1-\delta$. We show that, in this case, there exists a threshold $c_{k^*} \in C[h_t]$ such that types in $C^- = \{c \in C[h_t] : c < c_{k^*}\}$ accept with probability 1 and that types in $C^+ = \{c \in C[h_t] : c \geq c_{k^*}\}$ reject with probability 1. When C^- is nonempty, the principal offers transfer T_t in equation $(*)$ in the main text.

We start by showing that, at such a history (h_t, b_t), type $c_{\overline{k}[h_t]}$ takes the action with probability 0. Suppose to the contrary that type $c_{\overline{k}[h_t]}$ takes the action with positive probability, so that $T_t = c_{\overline{k}[h_t]}$. If this is so, then all types $c_i < c_{\overline{k}[h_t]}$ must take the action with probability 1. To see why, suppose this is not true, and let c_i be the highest type below $c_{\overline{k}[h_t]}$ that does not take the action with probability 1. Since equilibrium behavior under (σ, μ) coincides with equilibrium behavior under (σ^P, μ^P) at all times $\tau \geq t+1$, the payoff that type c_i obtains by rejecting the offer is at most $X(b_t, E_{\overline{k}[h_t]})(c_{\overline{k}[h_t]} - c_i)$. However, type c_i can guarantee herself a payoff of $(1-\delta+X(b_t, E_{\overline{k}[h_t]}))(c_{\overline{k}[h_t]} - c_i)$ by accepting the offer today and then mimicking type $c_{\overline{k}[h_t]}$ at all times $\tau \geq t+1$, a contradiction. Since $c_{\overline{k}[h_t]} < b_t$, then the principal would be strictly better off under an equilibrium in Σ^M_{k-1} that is identical to (σ, μ), except that at history (h_t, b_t), the principal makes offer $T_t = c_{\overline{k}[h_t]}$ that is rejected by type $c_{\overline{k}[h_t]}$ and accepted by all types $c_i < c_{\overline{k}[h_t]}$. This contradicts $(\sigma, \mu) \in \Sigma^M_K$. Hence, at history (h_t, b_t), type $c_{\overline{k}[h_t]}$ takes the action with probability 0.

Next we show that at history (h_t, b_t), there exists a threshold $c_{k^*} \in C[h_t]$ such that types in $C^- = \{c \in C[h_t] : c < c_{k^*}\}$ accept with probability 1 and that types in $C^+ = \{c \in C[h_t] : c \geq c_{k^*}\}$ reject with probability 1. The statement is true if all types reject the offer with probability 1. Suppose the set of types in $C[h_t]$ that accept the offer with positive probability is nonempty, and let $c_{j^*} < c_{\overline{k}[h_t]}$ be the highest type in this set. Since equilibrium behavior at times $\tau \geq t+1$ coincides with (σ^P, μ^P), type c_{j^*} obtains a payoff of $(1-\delta)(T_t - c_{j^*}) + \delta \times 0$. Let c_{k^*} be the lowest type in $\{c \in C[h_t] : c > c_{j^*}\}$. Note that the offer that the principal makes must satisfy $(*)$ in the main text:

$$(1-\delta)(T_t - c_{j^*}) = V^{(\sigma^P,\mu^P)}_{k^*}[h_t, b_t] + A^{(\sigma^P,\mu^P)}_{k^*}[h_t, b_t](c_{k^*} - c_{j^*}).$$

Indeed, this transfer leaves type c_{j^*} indifferent between accepting the offer and rejecting it. Since type c_{k^*} rejects the offer with probability 1, it must be that

$$V^{(\sigma^P,\mu^P)}_{k^*}[h_t, b_t] \geq (1-\delta)(T_t - c_{k^*}) \iff 1-\delta \geq A^{(\sigma^P,\mu^P)}_{k^*}[h_t, b_t]. \tag{SA.2}$$

We now show that type c_{j^*} accepts with probability 1. Indeed, the payoff that the principal obtains from type c_{j^*} from $t+1$ onward if this type accepts the offer is $(1-\delta)\mathbb{E}[\sum_{s=1}^{\infty} \delta^s \mathbf{1}_{b_{t+s} \in E_{j^*}}(b_{t+s} - c_{j^*})|b_t]$, which is the efficient payoff and is clearly higher than what she would obtain from this type if the type rejects the offer.[7]

[6]Indeed, starting from $t+1$, equilibrium behavior under (σ, μ) coincides with equilibrium behavior under (σ^P, μ^P). As a result, the profits that the principal obtains from each type of agent $c_i < c_{\overline{k}[h_t]}$ from $t+1$ onward do not depend on the relative likelihood that she assigns to type $c_{\overline{k}[h_t]}$. Moreover, the profits that she extracts from type $c_{\overline{k}[h_t]}$ from $t+1$ onward are the same regardless of whether this type accepts or not. These two observations imply that, at time t, the principal is better off making an offer that every type of agent rejects.

[7]Moreover, if some types $c_i < c_{j^*}$ were to reject the offer, the continuation payoff that the principal would get from them would be weakly higher if type c_{j^*} were to accept offer T_t with probability 1 than if type c_{j^*} were to reject the offer with positive probability. Indeed, if type c_{j^*} is not in the support of the principal's beliefs at time $t+1$, then types $c_i < c_{j^*}$ get smaller informational rents.

Next, we show that all types in $c_i \in C[h_t]$ with $c_i < c_{j^*}$ accept offer T_t with probability 1. Toward a contradiction, let c_i be the highest type below c_i that rejects T_t with positive probability. Since equilibrium behavior from $t+1$ onward under (σ, μ) coincides with equilibrium behavior under (σ^P, μ^P), type c_i obtains payoff $V_{k^*}^{(\sigma^P,\mu^P)}[h_t, b_t] + A_{k^*}^{(\sigma^P,\mu^P)}[h_t, b_t](c_{k^*} - c_i)$ from rejecting offer T_t. In contrast, the payoff that type c_i would obtain from accepting offer T_t and mimicking type c_{j^*} from time $t+1$ onward is $(1-\delta)(T_t - c_i) + X(b_t, E_{j^*})(c_{j^*} - c_i)$. Note that

$$\begin{aligned}
&(1-\delta)(T_t - c_i) + X(b_t, E_{j^*})(c_{j^*} - c_i) - V_{k^*}^{(\sigma^P,\mu^P)}[h_t, b_t] - A_{k^*}^{(\sigma^P,\mu^P)}[h_t, b_t](c_{k^*} - c_i) \\
&\quad = (c_{j^*} - c_i)\big(1 - \delta + X(b_t, E_{j^*}) - A_{k^*}^{(\sigma^P,\mu^P)}[h_t, b_t]\big) > 0,
\end{aligned}$$

where we used equation (SA.2).

The arguments above show that, at histories (h_t, b_t) with $X(b_t, E_{\overline{k}[h_t]}) \leq 1$, there exists a threshold $c_{k^*} \in C[h_t]$ such that types in $C^- = \{c \in C[h_t] : c < c_{k^*}\}$ accept with probability 1 and that types in $C^+ = \{c \in C[h_t] : c \geq c_{k^*}\}$ reject with probability 1. Since the threshold c_{k^*} is chosen optimally under equilibrium (σ^P, μ^P), under equilibrium (σ, μ) the principal would choose the same cutoff. Hence, at history (h_t, b_t), players' behavior under $(\sigma, \mu) \in \Sigma_K^M$ coincides with their behavior under (σ^P, μ^P). *Q.E.D.*

APPENDIX SA.3: FULL COMMITMENT

This appendix studies the problem of a principal who has full commitment power. For conciseness, we focus on the case in which there are two types of agents: $\mathcal{C} = \{c_1, c_2\}$, with $c_1 < c_2$. Let $\mu \in (0, 1)$ be the probability that the agent's cost is c_2.

The principal's problem is to choose processes $\{a_{i,t}, T_{i,t}\}$ for $i = 1, 2$, with $a_{i,t} \in \{0, 1\}$ and $T_{i,t} \in \mathbb{R}$, to solve

$$U^{\mathrm{FC}}(b) = \max_{\{a_{i,t}T_{i,t}\}_{i=1,2}} (1-\delta)\mathbb{E}\left[\sum_{t=0}^{\infty} \delta^t \big((1-\mu)(a_{1,t}b_t - T_{1,t}) + \mu(a_{2,t}b_t - T_{2,t})\big)\middle| b_0 = b\right] \tag{SA.3}$$

$$\text{subject to} \quad \mathbb{E}\left[\sum_{t=0}^{\infty} \delta^t (T_{i,t} - a_{i,t}c_i)\middle| b_0 = b\right] \geq 0 \quad \text{for } i = 1, 2 \quad \text{and}$$

$$\begin{aligned}
&\mathbb{E}\left[\sum_{t=0}^{\infty} \delta^t (T_{i,t} - a_{i,t}c_i)\middle| b_0 = b\right] \\
&\quad \geq \mathbb{E}\left[\sum_{t=0}^{\infty} \delta^t (T_{j,t} - a_{j,t}c_i)\middle| b_0 = b\right] \quad \text{for } i, j = 1, 2.
\end{aligned}$$

By familiar arguments, the participation constraint of type c_1 and the incentive compatibility constraint of type c_2 do not bind. The participation constraint of type c_2 and the incentive compatibility constraint of type c_1 hold with equality at the solution to (SA.3). Using these two constraints to solve for the expected discounted transfers

$(1-\delta)\mathbb{E}[\sum_{t=0}^{\infty}\delta^t T_{i,t}|b_0=b]$ for $i=1,2$ and replacing them into the objective yields

$$U^{\mathrm{FC}}(b)=\max_{\{a_{i,t}\}_{i=1,2}}(1-\delta)\mathbb{E}\Bigg[\sum_{t=0}^{\infty}\delta^t\bigg((1-\mu)a_{1,t}(b_t-c_1)+\mu a_{2,t}\bigg(b_t-c_2-\frac{(1-\mu)}{\mu}(c_2-c_1)\bigg)\bigg)\bigg|b_0=b\Bigg]. \tag{SA.4}$$

The solution to problem (SA.4) has $a_{1,t}=1$ if and only if (iff) $b_t\geq c_1$ (i.e., iff $b_t\in E_1$), and $a_{2,t}=1$ if and only if $b_t\geq c_2+\frac{(1-\mu)}{\mu}(c_2-c_1)>c_2$.

The following result shows that, in the presence of stochastic shocks, the principal's equilibrium payoffs can be close to her full commitment payoffs

PROPOSITION SA.1: *Let $\mathcal{C}=\{c_1,c_2\}$ and assume there exists $b\in E_2\setminus E_1$ with $X(b,E_2)=\varepsilon<1-\delta$. Then, at histories (h_t,b_t) with $C[h_t]=\{c_1,c_2\}$ and $b_t=b$,*

$$U^{\mathrm{FC}}(b_t)-U^{\sigma,\mu}[h_t,b_t]\leq(1-\mu)(c_2-c_1)\varepsilon.$$

PROOF: Note that at such a history, the principal can make a separating offer T with $(1-\delta)(T-c_1)=X(b,E_1)(c_2-c_1)$ that only low types accept. Conditional on the agent being a low type, the principal's profits are

$$(1-\delta)\mathbb{E}\Bigg[\sum_{\tau=t}^{\infty}\delta^{\tau-t}\mathbf{1}_{b_\tau\in E_1}(b_\tau-c_1)\bigg|b_t=b\Bigg]-X(b,E_1)(c_2-c_1).$$

Conditional on the agent's type being a high type, the principal's profits are

$$(1-\delta)\mathbb{E}\Bigg[\sum_{\tau=t}^{\infty}\delta^{\tau-t}\mathbf{1}_{b_\tau\in E_2}(b_\tau-c_2)\bigg|b_t=b\Bigg].$$

The principal's expected payoff at history (h_t,b_t) from making offer T is then

$$U^{\sigma,\mu}[h_t,b_t]=(1-\delta)\mathbb{E}\Bigg[\sum_{\tau=t}^{\infty}\delta^{\tau-t}\bigg((1-\mu)\mathbf{1}_{b_\tau\in E_1}(b_\tau-c_1)+\mu\mathbf{1}_{b_\tau\in E_2}\bigg(b_t-c_2-\frac{(1-\mu)}{\mu}(c_2-c_1)\bigg)\bigg)\bigg|b_t\Bigg]. \tag{SA.5}$$

The principal's full commitment payoffs are

$$U^{\mathrm{FC}}(b_t)=(1-\delta)\mathbb{E}\Bigg[\sum_{\tau=t}^{\infty}\delta^{\tau-t}\bigg((1-\mu)\mathbf{1}_{b_\tau\in E_1}(b_\tau-c_1)+\mu\mathbf{1}_{b_\tau\in \hat{E}_2}\bigg(b_t-c_2-\frac{(1-\mu)}{\mu}(c_2-c_1)\bigg)\bigg)\bigg|b_t\Bigg], \tag{SA.6}$$

where $\hat{E}_2 = \{b \in \mathcal{B} : b_t \geq c_2 + (1-\mu)(c_2 - c_1)/\mu\} \subset E_2$. Using (SA.5) and (SA.6),

$$\begin{aligned} U^{\mathrm{FC}}(b_t) - U^{\sigma,\mu}[h_t, b_t] &= -(1-\delta)\mu\mathbb{E}\left[\sum_{\tau=t}^{\infty} \delta^{\tau-t}\mathbf{1}_{b_\tau \in E_2 \setminus \hat{E}_2}\left(b_t - c_2 - \frac{(1-\mu)}{\mu}(c_2 - c_1)\right)\Big| b_t\right] \\ &\leq (1-\delta)\mu\mathbb{E}\left[\sum_{\tau=t}^{\infty} \delta^{\tau-t}\mathbf{1}_{b_\tau \in E_2 \setminus \hat{E}_2}\left(\frac{(1-\mu)}{\mu}(c_2 - c_1)\right)\Big| b_t\right] \\ &= (1-\mu)(c_2 - c_1)X(b_t, E_2 \setminus \hat{E}_2) \\ &\leq (1-\mu)(c_2 - c_1)\varepsilon, \end{aligned}$$

where the first inequality follows since $b_\tau \geq c_2$ for all $b_\tau \in E_2$, and the second inequality follows since $X(b_t, E_2 \setminus \hat{E}_2) \leq X(b_t, E_2) = \varepsilon$. *Q.E.D.*

APPENDIX SA.4: PATH DEPENDENCE WHEN SHOCKS ARE ERGODIC

In this appendix, we show by example that the equilibrium may exhibit long-run path dependence when the shock process is ergodic. Let $\mathcal{B} = \{b_L, b_{ML}, b_{MH}, b_H\}$, with $b_L < b_{ML} < b_{MH} < b_H$ and $\mathcal{C} = \{c_1, c_2, c_3\}$. Assume that the efficiency sets are $E_1 = E_2 = \{b_{ML}, b_{MH}, b_H\}$ and $E_3 = \{b_H\}$.

PROPOSITION SA.2: *Suppose that the transition matrix $[Q_{b,b'}]$ satisfies*

(a) $Q_{b,b'} > 0$ *for all* $b, b' \in \mathcal{B}$,

(b) $X(b_{MH}, \{b_H\}) > 1-\delta$, $X(b, \{b_H\}) < 1-\delta$ *for* $b = b_{ML}, b_L$, *and* $X(b_{ML}, \{b_{ML}\}) > 1-\delta$.

Then there exists $\varepsilon_1 > 0$, $\varepsilon_2 > 0$, $\Delta_1 > 0$, and $\Delta_2 > 0$ such that if $Q_{b,b_L} < \varepsilon_1$ for all $b \in \mathcal{B} \setminus \{b_L\}$ and $Q_{b,b_{ML}} < \varepsilon_2$ for all $b \in \mathcal{B} \setminus \{b_{ML}\}$, and if $|b_L - c_1| < \Delta_1$ and $|b_L - c_2| > \Delta_2$, the unique equilibrium satisfies the following properties:

(i) *For histories h_t such that $C[h_t] = \{c_1, c_2\}$, $\mu[h_{t'}] = \mu[h_t]$ for all $h_{t'} \succeq h_t$ (i.e., there is no more learning by the principal from time t onward).*

(ii) *For histories h_t such that $C[h_t] = \{c_2, c_3\}$, if $b_t = b_L$ or $b_t = b_{MH}$, types c_2 and c_3 take action $a = 0$; if $b_t = b_{ML}$, type c_2 takes action $a = 1$ and type c_3 takes action $a = 0$; and if $b_t = b_H$, types c_2 and c_3 take action $a = 1$.*

(iii) *For histories h_t such that $C[h_t] = \{c_1, c_2, c_3\}$, if $b_t = b_L$, type c_1 takes action $a = 1$ while types c_2 and c_3 take action $a = 0$; if $b_t = b_{ML}$, types c_1 and c_2 take action $a = 1$ and type c_3 takes action $a = 0$; if $b_t = b_{MH}$, all agent types take action $a = 0$; and if $b_t = b_H$, all agent types take action $a = 1$.*

We prove the three properties in Proposition SA.2 separately.

PROOF OF PROPERTY (I): Note first that, by Theorem 1, after such a history the principal makes a pooling offer $T = c_2$ that both types accept if $b_t \in E_2 = \{b_{ML}, b_{MH}, b_H\}$. To establish the result, we show that if $b_t = b_L$, types c_1 and c_2 take action $a = 0$ after history h_t. If the principal makes a separating offer that only a c_1 agent accepts, she pays a transfer $T_t = c_1 + \frac{1}{1-\delta}X(b_L, E_2)(c_2 - c_1)$ that compensates the low-cost agent for revealing his type. The principal's payoff from making such an offer, conditional on the agent being type c_1, is

$$\tilde{U}^{\text{sc}}[c_1] = (1-\delta)(b_L - T_t) + \mathbb{E}\left[\sum_{t'>t} \delta^{t'-t}(1-\delta)\mathbf{1}_{b_t \in E_1}(b_{t'} - c_1)\middle| b_t = b_L\right]$$
$$= (1-\delta)(b_L - c_1) + \sum_{b\in\{b_{ML}, b_{MH}, b_H\}} X\big(b_L, \{b\}\big)[b - c_2].$$

Her payoff from making that offer conditional on the agent's type being c_2 is $\tilde{U}^{\text{sc}}[c_2] = \sum_{b\in\{b_{ML}, b_{MH}, b_H\}} X(b_L, \{b\})[b - c_2]$. If she does not make a separating offer when $b_t = b_L$, she never learns the agent's type and gets a payoff $\tilde{U}^{\text{nsc}} = \sum_{b\in\{b_{ML}, b_{MH}, b_H\}} X(b_L, \{b\}) \times [b - c_2]$. Since $b_L - c_1 < 0$ by assumption, $\tilde{U}^{\text{nsc}} > \mu[h_t][c_1]\tilde{U}^{\text{sc}}[c_1] + \mu[h_t][c_2]\tilde{U}^{\text{sc}}[c_2]$ and, therefore, the principal does not make a separating offer. *Q.E.D.*

Proof of Property (ii): Theorem 1 implies that, after such a history, the principal makes a pooling offer $T = c_3$ that both types accept if $b_t \in E_3 = \{b_H\}$. Theorem 1 also implies that if $b_t = b_{MH}$, then after such a history, the principal makes an offer that both types reject (since $X(b_{MH}, \{b_H\}) > 1 - \delta$ by assumption). So it remains to show that after history h_t, the principal makes an offer that a c_2 agent accepts and a c_3 agent rejects if $b_t = b_{ML}$, and that the principal makes an offer that both types reject if $b_t = b_L$.

Suppose $b_t = b_{ML}$. Let $U[c_i]$ be the principal's value at history $(h_t, b_t = b_{ML})$ conditional on the agent's type being $c_i \in \{c_2, c_3\}$ and let V_i be the value of an agent of type c_i at history $(h_t, b_t = b_{ML})$. Note that $U[c_2] + V_2 \le (1-\delta)(b_{ML} - c_2) + \sum_{b\in\{b_{ML}, b_{MH}, b_H\}} X(b_{ML}, \{b\})[b - c_2]$, since the right-hand side of this equation corresponds to the efficient total payoff when the agent is of type c_2 (i.e., the agent taking the action if and only if the state is in E_2.) Note also that incentive compatibility implies $V_2 \ge X(b_{ML}, \{b_H\})(c_2 - c_3)$, since a c_2 agent can mimic a c_3 agent forever and obtain $X(b_{ML}, \{b_H\})(c_2 - c_3)$. It thus follows that $U[c_2] \le (1-\delta)(b_{ML} - c_2) + X(b_{ML}, \{b_H\}) \times [b_H - c_3] + \sum_{s\in\{b_{ML}, b_{MH}\}} X(b_{ML}, \{b\})[b - c_2]$.

When $b_t = b_{ML}$, if the principal makes an offer that only a c_2 agent accepts, the offer must satisfy $T_t = c_2 + \frac{1}{1-\delta}X(b_{ML}, \{b_H\})(c_3 - c_2) < c_3$. The principal's payoff from making such an offer when the agent's type is c_2 is

$$\overline{U}[c_2] = (1-\delta)(b_{ML} - T_t) + \sum_{b\in\{b_{ML}, b_{MH}, b_H\}} X\big(b_{ML}, \{b\}\big)[b - c_2]$$
$$= (1-\delta)(b_{ML} - c_2) + X\big(b_{ML}, \{b_H\}\big)[b_H - c_3] + \sum_{b\in\{b_{ML}, b_{MH}\}} X\big(b_{ML}, \{b\}\big)[b - c_2],$$

which, from the arguments in the previous paragraph, is the highest payoff that the principal can ever get from a c_2 agent after history $(h_t, b_t = b_{ML})$. Hence, it is optimal for the principal to make such a separating offer.[8]

Suppose next that $b_t = b_L$. If the principal makes an offer that a c_2 agent accepts and a c_3 agent rejects, she pays a transfer $T_t = c_2 + \frac{1}{1-\delta}X(b_L, E_3)(c_3 - c_2)$. Thus, the principal's

[8]Indeed, the principal's payoff from making an offer equal to T_t when the agent's type is c_3 is $X(2, \{4\})[b(4) - c_3]$, which is also the most that she can extract from an agent of type c_3.

payoff from making such an offer, conditional on the agent being type c_2, is

$$\begin{aligned}\tilde{U}^{\text{sc}}[c_2] &= (1-\delta)(b_L - T_t) + \sum_{b\in\{b_{ML},b_{MH},b_H\}} X\big(b_L,\{b\}\big)[b-c_2] \\ &= (1-\delta)(b_L - c_2) + X\big(b_L,\{b_H\}\big)[b_H - c_3] + \sum_{b\in\{b_{ML},b_{MH}\}} X\big(b_L,\{b\}\big)[b-c_2].\end{aligned}$$

If the principal makes an offer that both types reject when $b_t = b_L$, then by the arguments above she learns the agent's type the first time at which shock b_{ML} is reached. Let $\check{t}$ be the random variable that indicates the next date at which shock b_{ML} is realized. Then, conditional on the agent's type being c_2, the principal's payoff from making an offer that both types reject when $b_t = b_L$ is

$$\begin{aligned}\tilde{U}^{\text{nsc}}[c_2] &= (1-\delta)\mathbb{E}\left[\sum_{t'=t+1}^{\check{t}-1} \delta^{t'-t}\mathbf{1}_{b_{t'}=b_H}(b_H - c_3)\Big| b_t = b_L\right] \\ &\quad + \mathbb{E}\left[\delta^{\check{t}-t}\left((1-\delta)(b_{ML} - T_{\check{t}}) + \sum_{b\in\{b_{ML},b_{MH},b_H\}} X\big(b_{ML},\{b\}\big)[b-c_2]\right)\Big| b_t = b_L\right].\end{aligned}$$

The offer $T_{\check{t}}$ that the principal makes at time $\check{t}$ satisfies $T_{\check{t}} = c_2 + \frac{1}{1-\delta}X(b_{ML},\{b_H\})(c_3 - c_2)$. Using this in the equation above yields

$$\begin{aligned}\tilde{U}^{\text{nsc}}[c_2] &= X\big(b_L,\{b_H\}\big)[b_H - c_3] + X\big(b_L,\{b_{ML}\}\big)[b_{ML} - c_2] \\ &\quad + \mathbb{E}\big[\delta^{\check{t}-t}|b_t = b_L\big]X\big(b_{ML},\{b_{MH}\}\big)[b_{MH} - c_2].\end{aligned}$$

Then we have

$$\begin{aligned}\tilde{U}^{\text{nsc}}[c_2] - \tilde{U}^{\text{sc}}[c_2] &= -(1-\delta)[b_L - c_2] - \big[X\big(b_L,\{b_{MH}\}\big) \\ &\quad - \mathbb{E}\big[\delta^{\check{t}-t}|b_t = b_L\big]X\big(b_{ML},\{b_{MH}\}\big)\big][b_{MH} - c_2].\end{aligned}$$

Since $b_L < c_2$ by assumption, there exists $\Delta_2^1 > 0$ such that if $(1-\delta)(c_2 - b_L) > \Delta_2^1$, the expression above is positive. Since the principal's payoff conditional on the agent's type being c_3 is the same regardless of whether she makes a separating offer or not when $b_t = b_L$ (i.e., in either case the principal earns $X(b_L,\{b_H\})(b_H - c3)$), when this condition holds the principal does not make an offer that c_2 accepts and c_3 rejects when $b_t = b_L$. *Q.E.D.*

PROOF OF PROPERTY (III): Suppose $C[h_t] = \{c_1, c_2, c_3\}$. Theorem 1 implies that all agent types take action $a = 1$ if $b_t = b_H$ and all agent types take action $a = 0$ if $b_t = b_{MH}$ (this last claim follows since $X(b_{MH},\{b_H\}) > 1-\delta$).

Suppose next that $C[h_t] = \{c_1, c_2, c_3\}$ and $b_t = b_{ML}$. Note that, by Lemma A.1, an agent with type c_3 takes action $a = 0$ if $b_t = b_{ML} \notin E_3 = \{b_H\}$. We first claim that if the principal makes an offer that only a subset of types accept at state b_{ML}, then this offer must be such that types in $\{c_1, c_2\}$ take action $a = 1$ and type c_3 takes action $a = 0$. To see this, suppose that she instead makes an offer that only an agent with type c_1 accepts and that agents with types in $\{c_2, c_3\}$ reject. The offer that she makes in this case satisfies $(1-\delta)(T_t - c_1) = V_2^{(\sigma,\mu)}[h_t, b_t] + A_2^{(\sigma,\mu)}[h_t, b_t](c_2 - c_1)$. By property (ii) above, under this proposed equilibrium, a c_2 agent will from period $t+1$ onward take the action at all times $t' > t$ such that $b_{t'} = b_{ML}$.[9] Therefore, $A_2^{(\sigma,\mu)}[h_t, b_t] \geq X(b_{ML}, \{b_{ML}\}) > 1-\delta$, where the last inequality follows by assumption. The payoff that an agent of type c_2 obtains by accepting offer T_t at time t is bounded below by $(1-\delta)(T_t - c_2) = (1-\delta)(c_1 - c_2) + V_2^{(\sigma,\mu)}[h_t, b_t] + A_2^{(\sigma,\mu)}[h_t, b_t](c_2 - c_1) > V_2^{(\sigma,\mu)}[h_t, b_t]$, where the inequality follows since $A_2^{(\sigma,\mu)}[h_t, b_t] > 1-\delta$. Thus, type c_2 strictly prefers to accept the offer, a contradiction. Therefore, when $C[h_t] = \{c_1, c_2, c_3\}$ and $b_t = b_{ML}$, either the principal makes an offer that only types in $\{c_1, c_2\}$ accept or she makes an offer that all types reject.

We now show that under the conditions in the lemma, the principal makes an offer that types in $\{c_1, c_2\}$ accept and type c_3 rejects when $b_t = b_{ML}$ and $C[h_t] = \{c_1, c_2, c_3\}$. If she makes an offer that agents with cost in $\{c_1, c_2\}$ accept and a c_3 agent rejects, then she pays a transfer $T_t = c_2 + \frac{1}{1-\delta} X(b_{ML}, \{b_H\})(c_3 - c_2)$. Note then that, by property (i) above, when the agent's cost is in $\{c_1, c_2\}$, the principal stops learning: for all times $t' > t$, the principal makes an offer $T_{t'} = c_2$ that both types accept when $b_{t'} \in E_2$, and she makes a low offer $T_{t'} = 0$ that both types reject when $b_{t'} \notin E_2$. Therefore, conditional on the agent's type being either c_1 or c_2, the principal's payoff from making, at time t, an offer T_t that agents with cost in $\{c_1, c_2\}$ accept and a c_3 agent rejects is

$$\begin{aligned}
\hat{U}^{\mathrm{sc}}\big[\{c_1, c_2\}\big] &= (1-\delta)(b_{ML} - T_t) + \sum_{b \in \{b_{ML}, b_{MH}, b_H\}} X\big(b_{ML}, \{b\}\big)[b - c_2] \\
&= (1-\delta)(b_{ML} - c_2) + X\big(b_{ML}, \{b_H\}\big)[b_H - c_3] \\
&\quad + \sum_{b \in \{b_{ML}, b_{MH}\}} X\big(b_{ML}, \{b\}\big)[b - c_2].
\end{aligned}$$

On the other hand, if she does not make an offer that a subset of types accept when $b_t = b_{ML}$, then the principal's payoffs conditional on the agent being of type $c_i \in \{c_1, c_2\}$ are bounded above by

$$\hat{U}^{\mathrm{nsc}}[c_i] = \mathbb{E}\left[\sum_{t'=t}^{\hat{t}-1} \delta^{t'-t}(1-\delta)\mathbf{1}_{b_{t'}=b_H}(b_H - c_3) + \delta^{\hat{t}-t} \sum_{b \in E_i} X\big(b_L, \{b\}\big)(b - c_i) | b_t = b_{ML}\right],$$

where $\hat{t}$ denotes the next period that state b_L is realized.[10] Note that there exists $\varepsilon_1 > 0$ small such that if $Q_{b,b_L} < \varepsilon_1$ for all $b \neq b_L$, then $\hat{U}^{\mathrm{sc}}[\{c_1, c_2\}] > \hat{U}^{\mathrm{nsc}}[c_i]$ for $i = 1, 2$. Finally, note that the payoff that the principal obtains from an agent of type c_3 at history h_t when $b_t = b_{ML}$ is $X(b_{ML}, \{b_H\})(b_H - c_3)$, regardless of the principal's offer. Therefore, if $Q_{b,b_L} < \varepsilon_1$ for all $b \neq b_L$, when $C[h_t] = \{c_1, c_2, c_3\}$ and $b_t = b_{ML}$, the principal makes an offer T_t that only types in $\{c_1, c_2\}$ accept.

[9]Under the proposed equilibrium, if the offer is rejected, the principal learns that the agent's type is in $\{c_2, c_3\}$. By property (ii), if the agent's type is c_2, the principal will learn the agent's type the next time the shock is b_{ML} (because at that time type c_2 takes the action, while type c_3 does not), and from that point onward the agent will take the action when the shock is in $E_2 = \{b_{ML}, b_{MH}, b_H\}$.

[10]To see why, note that if no type of agent takes the productive action when $C[h_t] = \{c_1, c_2, c_3\}$ and $b_t = b_{ML}$, then the principal can only learn the agent's type when state b_L is realized (i.e., at time $\hat{t}$). At times before $\hat{t}$, all agent types take the action if the shock is b_H (and the principal pays transfer $T = c_3$) and no agent type takes the action at states b_{ML} or b_{MH}. After time $\hat{t}$, the payoff that the principal gets from type c_i is bounded above by her first-best payoff $\sum_{b \in E_i} X(b_L, \{b\})(b - c_i)$.

Finally, we show that when $C[h_t] = \{c_1, c_2, c_3\}$ and $b_t = b_L$, the principal makes an offer that only type c_1 accepts. Let $\check{t}$ be the random variable that indicates the next date at which state b_{ML} is realized. If the principal makes an offer T_t that only a c_1 agent accepts, this offer satisfies

$$\begin{aligned}(1-\delta)(T_t - c_1) &= V_2^{(\sigma,\mu)}[h_t, b_L] + A_2^{(\sigma,\mu)}[h_t, b_L](c_2 - c_1) \\ &= X\big(b_L, \{b_H\}\big)(c_3 - c_1) \\ &\quad + \big[X\big(b_L, \{b_{ML}\}\big) + \mathbb{E}\big[\delta^{\check{t}-t} | b_t = b_L\big] X\big(b_{ML}, \{b_{MH}\}\big)\big](c_2 - c_1),\end{aligned} \tag{SA.7}$$

where the second equality follows since $V_2^{(\sigma,\mu)}[h_t, b_L] = A_3^{(\sigma,\mu)}[h_t, b_L](c_3 - c_2) = X(b_L, \{b_H\})(c_3 - c_2)$ and since, by property (ii), when the support of the principal's beliefs is $\{c_2, c_3\}$ and the agent's type is c_2, the principal learns the agent's type at time $\check{t}$.[11] Therefore, conditional on the agent's type being c_1, the principal's equilibrium payoff from making an offer that only an agent with cost c_1 accepts at state b_L is

$$\begin{aligned}\check{U}^{\mathrm{sc}}[c_1] &= (1-\delta)(b_L - T_t) + \sum_{b \in \{b_{ML}, b_{MH}, b_H\}} X\big(b_L, \{b\}\big)[b - c_1] \\ &= (1-\delta)(b_L - c_1) + X\big(b_L, \{b_H\}\big)[b_H - c_3] + X\big(b_L, \{b_{MH}\}\big)[b_{MH} - c_1] \\ &\quad + X\big(b_L, \{b_{ML}\}\big)[b_{ML} - c_2] - \mathbb{E}\big[\delta^{\check{t}-t} | b_t = b_L\big] X\big(b_{ML}, \{b_{MH}\}\big)(c_2 - c_1),\end{aligned}$$

where the second line follows from substituting the transfer in (SA.7). On the other hand, the principal's payoff from making such an offer at state b_L, conditional on the agent's type being c_2, is

$$\begin{aligned}\check{U}^{\mathrm{sc}}[c_2] &= (1-\delta)\mathbb{E}\left[\sum_{t'=t}^{\check{t}-1} \delta^{t'-t} \mathbf{1}_{b_{t'} = b_H}(b_H - c_3) \middle| b_t = b_L\right] \\ &\quad + (1-\delta)\mathbb{E}\left[\delta^{\check{t}-t}\left((b_{ML} - c_2) - \frac{X\big(b_{ML}, \{b_H\}\big)(c_3 - c_2)}{1-\delta}\right)\right. \\ &\quad \left. + \sum_{t'=\check{t}+1}^{\infty} \delta^{t'-t} \mathbf{1}_{b_{t'} \in E_2}(b_{t'} - c_2) \middle| b_t = b_L\right] \\ &= X\big(b_L, \{b_H\}\big)(b_H - c_3) + X\big(b_L, \{b_{ML}\}\big)(b_{ML} - c_2) \\ &\quad + \mathbb{E}\big[\delta^{\check{t}-t} X\big(b_{ML}, \{b_{MH}\}\big) | b_t = b_L\big](b_{MH} - c_2),\end{aligned}$$

[11]The fact that the principal learns the agent's type at time $\check{t}$ implies that

$$\begin{aligned}A_2^{(\sigma,\mu)}[h_t, b_L] &= (1-\delta)\mathbb{E}\left[\sum_{t'=t}^{\check{t}-1} \delta^{t'-t} \mathbf{1}_{b_{t'} = b_H} + \delta^{\check{t}-t} \sum_{t'=\check{t}}^{\infty} \delta^{t'-\check{t}} \mathbf{1}_{b_{t'} \in E_2} \middle| b_t = b_L\right] \\ &= X\big(b_L, \{b_H\}\big) + X\big(b_L, \{b_{ML}\}\big) + \mathbb{E}\big[\delta^{\check{t}-t} X\big(b_{ML}, \{b_{MH}\}\big) | b_t = b_L\big].\end{aligned}$$

where we used the fact that when the support of her beliefs is $\{c_2, c_3\}$, the principal makes an offer that only a c_2 agent accepts when the state is b_{ML} (the offer that she makes at that point is $T = c_2 + \frac{1}{1-\delta} X(b_{ML}, \{b_H\})(c_3 - c_2)$).

Alternatively, suppose the principal makes an offer that both c_1 and c_2 accept but c_3 rejects. Then she pays a transfer $T_t = c_2 + \frac{1}{1-\delta} X(b_L, \{b_H\})(c_3 - c_2)$; thus, her payoff from learning that the agent's type is in $\{c_1, c_2\}$ in state b_L is

$$\begin{aligned}\bar{U}^{\text{sc}}\big[\{c_1, c_2\}\big] &= (1-\delta)(b_L - T_t) + \sum_{b \in \{b_{ML}, b_{MH}, b_H\}} X\big(b_L, \{b\}\big)(b - c_2) \\ &= (1-\delta)(b_L - c_2) + X\big(b_L, \{b_H\}\big)[b_H - c_3] \\ &\quad + X\big(b_L, \{b_{ML}\}\big)[b_{ML} - c_2] + X\big(b_L, \{b_{MH}\}\big)[b_{MH} - c_2],\end{aligned}$$

where we used the fact that the principal never learns anything more about the agent's type when the support of her beliefs is $\{c_1, c_2\}$ (see property (i) above). Note that there exists $\varepsilon_2 > 0$ and $\Delta_2^2 > 0$ such that if $Q_{b,b_{ML}} < \varepsilon_2$ for all $b \neq b_{ML}$ and if $c_2 - b_L > \Delta_2 = \max\{\Delta_2^1, \Delta_2^2\}$, then the following two inequalities hold:

$$\begin{aligned}\check{U}^{\text{sc}}[c_1] - \bar{U}^{\text{sc}}\big[\{c_1, c_2\}\big] &= \big[1 - \delta + X\big(b_L, \{b_{MH}\}\big) - \mathbb{E}\big[\delta^{\check{t}-t} | b_t = b_L\big] X\big(b_{ML}, \{b_{MH}\}\big)\big] \\ &\quad \times (c_2 - c_1) \\ &> 0 \\ \check{U}^{\text{sc}}[c_2] - \bar{U}^{\text{sc}}\big[\{c_1, c_2\}\big] &= \big[E\big[\delta^{\check{t}-t} X\big(b_{ML}, \{b_{MH}\}\big) | b_t = b_L\big] - X\big(b_L, \{b_{MH}\}\big)\big](b_{MH} - c_2) \\ &\quad - (1-\delta)(b_L - c_2) \\ &> 0.\end{aligned}$$

Therefore, under these conditions, at state b_L the principal strictly prefers to make an offer that a c_1 agent accepts and agents with cost $c \in \{c_2, c_3\}$ reject than to make an offer that agents with cost in $\{c_1, c_2\}$ accept and a c_3 agent rejects.

However, the principal may choose to make an offer that all agent types reject when $b_t = b_L$ and $C[h_t] = \{c_1, c_2, c_3\}$. In this case, by the arguments above, the next time the state is equal to b_{ML}, the principal will make an offer that only types in $\{c_1, c_2\}$ accept. The offer that she makes in this case is such that $(1-\delta)(T - c_2) = X(b_{ML}, \{b_H\})(c_3 - c_2)$. Then, from that point onward, she will never learn more (by property (i) above). In this case, the principal's payoff conditional on the agent's type being $\{c_1, c_2\}$ is

$$\begin{aligned}\bar{U}^{\text{nsc}} &= (1-\delta)\mathbb{E}\left[\sum_{\tau=t}^{\check{t}-1} \mathbf{1}_{b_\tau = b_H}(b_\tau - c_3) \middle| b_t = b_L\right] \\ &\quad + \mathbb{E}\left[\delta^{\check{t}-t}(1-\delta)(b_{ML} - T) + \sum_{b \in E_2} X\big(b_{ML}, \{b\}\big)(b - c_2) \middle| b_t = b_L\right] \\ &= X\big(b_L, \{b_H\}\big)[b_H - c_3] + X\big(b_L, \{b_{ML}\}\big)[b_{ML} - c_2] \\ &\quad + \mathbb{E}\big[\delta^{\check{t}-t} | b_t = b_L\big] X\big(b_{ML}, \{b_{MH}\}\big)[b_{MH} - c_2],\end{aligned}$$

where $\check{t}$ is the random variable that indicates the next date at which state b_{ML} is realized. Note that there exists $\varepsilon_2 > 0$ and $\Delta_1 > 0$ such that if $Q_{b,b_{ML}} < \varepsilon_2$ for all $b \neq b_{ML}$ and if

$b_L - c_1 > -\Delta_1$, then the following equations hold:

$$\begin{aligned}\check{U}^{\text{sc}}[c_1] - \bar{U}^{\text{nsc}} &= (1-\delta)(b_L - c_1) \\ &\quad + \big[X\big(b_L, \{b_{MH}\}\big) - \mathbb{E}\big[\delta^{\check{t}-t} | b_t = b_L\big] X\big(b_{ML}, \{b_{MH}\}\big)\big][b_{MH} - c_1] \\ &> 0, \\ \check{U}^{\text{sc}}[c_2] - \bar{U}^{\text{nsc}} &= 0.\end{aligned}$$

Therefore, under these conditions, the principal makes an offer that type c_1 accepts and types in $\{c_2, c_3\}$ reject when $C[h_t] = \{c_1, c_2, c_3\}$ and $b_t = b_L$. *Q.E.D.*

Properties (i)–(iii) in Proposition SA.2 imply that the equilibrium exhibits long-run path dependence. Suppose that the agent's type is c_1. Then properties (i)–(iii) imply that the principal eventually learns the agent's type if and only if $t(b_L) := \min\{t \geq 0 : b_t = b_L\} < t(b_{ML}) := \min\{t \geq 0 : b_t = b_{ML}\}$ (i.e., if state b_L is visited before state b_{ML}). Indeed, if b_L is visited before b_{ML}, at time $t(b_L)$, the principal will learn that the agent's type is c_1 (see property (iii)). From that point onward, the agent will take the productive action at all periods $t > t(b_L)$ such that $b_t \in E_1$ at cost c_1 for the principal.

In contrast, if b_{ML} is visited before b_L, at time $t(b_{ML})$, the principal will learn that the agent's type is in $\{c_1, c_2\}$ (see property (iii)). From that point onward there will be no more learning (property (i)). As a consequence, the agent will take the productive action at all periods $t > t(b_{ML})$ such that $b_t \in E_2 = E_1$ at cost c_2 for the principal (this follows from Theorem 1(i)).

REFERENCE

PESKI, M. (2008): "Repeated Games With Incomplete Information on One Side," *Theoretical Economics*, 3 (1), 29–84. [3]

Co-editor Dirk Bergemann handled this manuscript.

Manuscript received 22 September, 2016; final version accepted 19 August, 2017; available online 21 August, 2017.