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

PROGRESSIVE LEARNING

AVIDIT ACHARYA
Department of Political Science, Stanford University

JUAN ORTNER
Department of Economics, Boston University

We study a dynamic principal–agent relationship with adverse selection and limited commitment. We show that when the relationship is subject to productivity shocks, the principal may be able to improve her value over time by progressively learning the agent's private information. She may even achieve her first-best payoff in the long run. The relationship may also exhibit path dependence, with early shocks determining the principal's long-run value. These findings contrast sharply with the results of the ratchet effect literature, in which the principal persistently obtains low payoffs, giving up substantial informational rents to the agent.

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

1. INTRODUCTION

CONSIDER A LONG-TERM RELATIONSHIP between an agent who has persistent private information and a principal who cannot commit to long-term contracts. If the parties are sufficiently forward-looking, then the relationship is subject to the ratchet effect: the agent is unwilling to disclose his private information, fearing that the principal will update the terms of his contract. This limits the principal's ability to learn the agent's private information and reduces her value from the relationship.

The ratchet effect literature has shed light on many economic applications including planning problems (Freixas, Guesnerie, and Tirole (1985)), labor contracting (Gibbons (1987), Dewatripont (1989)), regulation (Laffont and Tirole (1988)), optimal taxation (Dillen and Lundholm (1996)), repeated buyer–seller relationships (Hart and Tirole (1988), Schmidt (1993)), and relational contracting (Halac (2012), Malcomson (2016)).

A natural feature in virtually all of these applications is that productivity shocks to the economy have the potential to change the incentive environment over time. In this paper, we show that the classic ratchet effect results may not hold when the principal–agent relationship is subject to time-varying productivity shocks. In particular, the principal may gradually learn the agent's private information, which increases the value that she obtains from the relationship over time. The principal may even achieve her first-best payoff in the long run.

We study a stochastic game played between a principal and an agent. At each period, the principal offers the agent a transfer in exchange for taking an action that benefits her. The principal is able to observe the agent's action, but the agent's cost of taking the action

Avidit Acharya: avidit@stanford.edu
Juan Ortner: jortner@bu.edu
For helpful comments, we would like to thank Steve Callander, Bob Gibbons, Marina Halac, Bart Lipman, John Patty, Alan Wiseman, Stephane Wolton, and seminar audiences at Berkeley, Boston University, Collegio Carlo Alberto, the University of Hong Kong Summer Microeconomics Seminar, the LSE/NYU political economy conference, the NBER Org Econ working group, Northwestern Kellogg Strategy, Princeton, and Stanford. The paper also benefited from the thoughtful feedback of a co-editor and five anonymous referees.

 DOI: 10.3982/ECTA14718

is his private information and is constant over time. The principal has short-term, but not long-term, commitment power: she can credibly promise to pay a transfer in the current period if the agent takes the action, but cannot commit to future transfers. The realization of a productivity shock affects the size of the benefit that the principal obtains from having the agent take the action. The realization of the current period shock is publicly observed by both the principal and the agent at the start of the period, and the shock evolves over time as a Markov process.

Hart and Tirole (1988) and Schmidt (1993) study the special case of our model in which productivity is constant over time. We show how the equilibrium of this special case differs qualitatively from the equilibrium of our model in which productivity changes over time. The three main differences are as follows.

First, we find that in the presence of productivity shocks, the equilibrium may be persistently inefficient. This contrasts with the equilibrium of the model without the shocks, which is efficient.

Second, productivity shocks give the principal the opportunity to progressively learn the agent's private information. As a result, the principal's value from the relationship gradually improves over time. We show that under natural assumptions, the principal is only able to get the agent to disclose some of his private information when productivity is low; that is, learning takes place in "bad times." We also show that productivity shocks may enable the principal to obtain profits that are arbitrarily close to her full commitment profits. Last, we derive conditions under which the principal ends up fully learning the agent's private information and attains her first-best payoffs in the long run.

Third, we show that learning by the principal may be path dependent: the degree to which the principal learns the agent's private information may depend critically on the order in which productivity shocks were realized early on in the relationship. This is true even when the process governing the evolution of productivity is ergodic. As a result, early shocks can have a lasting impact on the principal's value from the relationship.

Our model generates two testable predictions. First, the agent's performance will typically be higher after the realization of negative productivity shocks. This is consistent with Lazear, Shaw, and Stanton (2016), who find evidence that workers' productivity increases following a recession. Second, there will be hysteresis in the agent's compensation: the current wage of the agent is negatively affected by previous negative shocks. This resonates with Kahn (2010) and Oreopoulos, Von Wachter, and Heisz (2012), who find evidence that recessions have a long lasting impact on workers' compensation.

The key feature of our model that drives these dynamics is that the agent's incentive to conceal his private information changes over time. When current productivity is low and the future looks dim, the informational rents that low-cost types expect to earn by mimicking a higher-cost type are small. When these rents are small, it is cheap for the principal to get a low-cost agent to reveal his private information. These changes in the cost of inducing information disclosure make it possible for the principal to progressively screen the different types of agents, giving rise to our equilibrium dynamics.

Related Literature

Our work relates to prior papers that have suggested different ways to mitigate the ratchet effect. Kanemoto and MacLeod (1992) show that competition for second-hand workers may alleviate the ratchet effect. Carmichael and MacLeod (2000) show that the threat of future punishment may deter the principal from updating the terms of the agent's contract, mitigating the ratchet effect. Fiocco and Strausz (2015) show that the principal can incentivize information disclosure by delegating contracting to an independent third party. Our paper differs from these studies in that we do not introduce external sources of contract enforcement; neither do we reintroduce commitment by allowing for non-Markovian strategies.

Instead, we focus on the role that shocks play in ameliorating the principal's commitment problem. This connects our paper with Ortner (2017), who considers a durable goods monopolist who lacks commitment power and who faces time-varying production costs. In contrast to the classic results on the Coase conjecture (Fudenberg, Levine, and Tirole (1985), Gul, Sonnenschein, and Wilson (1986)), Ortner (2017) shows that time-varying costs may enable the monopolist to extract rents from high value buyers. A key difference between Ortner (2017) and the current paper is that the interaction between the monopolist and the buyers is one-shot in the Coasian environment. As a result, issues of information revelation, which are central to the current paper, are absent in that model.[1]

Blume (1998) generalizes the Hart and Tirole (1988) model to a setting in which the consumer's valuation has both permanent and transient components. Blume (1998) shows that optimal renegotiation-proof contracts in this environment give the buyer the chance to exit in the future in case his valuation falls. Gerardi and Maestri (2015) study a dynamic contracting model with no commitment in which the agent's private information affects his marginal cost of effort. They find that the principal's lack of commitment may lead her to offer inefficient pooling contracts.

Our model is strategically equivalent to a setting in which the agent has at each period an outside option, whose value varies over time and is publicly observed. This relates our model to papers studying how outside options affect equilibrium dynamics in the classic Coasian model (Fuchs and Skrzypacz (2010), Board and Pycia (2014), Hwang and Li (2017)).[2] The key difference, again, is that we study the effect that time-varying outside options have in settings with repeated interaction.[3]

The path-dependence result relates our paper to a series of recent studies in organization economics that attempt to explain the persistent performance differences among seemingly identical firms (Gibbons and Henderson (2012)). Chassang (2010) shows that path dependence may arise when a principal must learn how to effectively monitor the agent. Li and Matouschek (2013) study relational contracting environments in which the principal has private information, and they show that this private information may give rise to cycles. Callander and Matouschek (2014) show that persistent performance differences may arise when managers engage in trial and error experimentation. Halac and Prat (2016) show that path dependence arises due to the agent's changing beliefs about the principal's monitoring ability. We add to this literature by providing a new explanation for persistent performance differences, with new testable implications. Our results imply that firms that experience negative shocks earlier will later be more productive.

[1]The current paper also differs from Ortner (2017) in terms of results. Ortner (2017) shows that the monopolist's ability to extract rents diminishes as the support of the distribution of consumer values becomes dense. In contrast, the equilibrium dynamics of our model hold independently of how dense the support of the agent's cost distribution is.

[2]See also Compte and Jehiel (2002), who study the effect that outside options have in models of reputational bargaining.

[3]Our model also relates to Kennan (2001), who studies a bilateral bargaining game in which a long-run seller faces a long-run buyer. The buyer is privately informed about her valuation, which evolves over time as a Markov chain. Kennan (2001) shows that time-varying private information gives rise to cycles in which the seller's offer depends on the buyer's past purchasing decisions.

Finally, our paper relates to a broader literature on dynamic games with private information (Hart (1985), Sorin (1999), Wiseman (2005), Peski (2008), Peski (2014)). In this literature our paper relates closely to work by Watson (1999, 2002), who studies a private information partnership game and shows that the value of the partnership increases over time as the players gradually increase the stakes of their relationship to screen out bad types.

2. TWO-PERIOD EXAMPLE

Consider the following two-period game played between a principal and an agent. At $t=0$, the agent learns her cost of work $c \in \{c_L, c_H\}$. Let $\mu \in (0,1)$ be the probability that the agent's cost is c_L. At the start of each period $t=0,1$, the principal's benefit $b_t \in \{b_L, b_H\}$ from having the agent work is publicly revealed. After observing b_t, the principal offers the agent a transfer $T_t \geq 0$ for working. The agent then publicly chooses whether or not to work. The payoffs of the principal and an agent of type c are

$$(1-\delta)(b_0 - T_0)a_0 + \delta(b_1 - T_1)a_1,$$
$$(1-\delta)(T_0 - c)a_0 + \delta(T_1 - c)a_1,$$

where $a_t \in \{0,1\}$ denotes whether or not the agent works in period $t=0,1$ and $\delta \in (0,1)$ measures the importance of period $t=1$ relative to period $t=0$. We assume

$$0 \leq c_L < b_L < c_H < b_H \quad \text{and} \quad \mu < \frac{b_H - c_H}{b_H - c_L} =: \overline{\mu}.$$

We further assume that the benefit b_t is independently and identically distributed (i.i.d.) over time, with $\text{prob}(b_t = b_L) = q \in [0,1]$ for $t=0,1$. We consider pure strategy equilibria of this game.

Consider play at $t=1$. Since we focus on pure strategy equilibria, on path at the start of $t=1$, the principal' beliefs are equal to her prior or are degenerate. If the principal's beliefs are equal to her prior, she finds it optimal to offer a transfer $T_1 = c_H$ that both types accept if $b_1 = b_H$ (since $\mu < \overline{\mu}$), and she finds it optimal to offer transfer $T_1 = c_L$ that only a low-cost type accepts if $b_1 = b_L$. If the principal learned that the agent's cost is c, she finds it optimal to offer $T_1 = c$, which the agent accepts if and only if $b_1 > c$.

Consider now play at $t=0$. Suppose first that $b_0 = b_L$. In this case, the principal must choose between two options: make a low offer that both types reject or make a higher offer that only the low-cost type accepts. Making an offer that both types accept is not profitable since $b_L < c_H$. Suppose the principal makes a separating offer T_0 that only a low-cost type accepts. Note that a low-cost agent reveals his private information by accepting, so his payoff is $(1-\delta)(T_0 - c_L) + \delta 0$. Also note that the low-cost type can obtain a payoff of $\delta(1-q)(c_H - c_L)$ by rejecting the offer, so we must have $T_0 \geq c_L + \frac{\delta}{1-\delta}(1-q)(c_H - c_L)$. Since the high-cost type rejects offer T_0 if and only if $T_0 \leq c_H$, we must have $c_H \geq c_L + \frac{\delta}{1-\delta}(1-q)(c_H - c_L)$ or

$$\frac{\delta}{1-\delta}(1-q) \leq 1. \tag{1}$$

When the future is sufficiently valuable (i.e., $\delta > 1/2$), this inequality holds only if the probability $1-q$ of high productivity tomorrow is low enough; that is, if the future looks dim. When (1) holds, the principal finds it optimal to make a separating offer, since such an offer gets the low-cost type to work at time $t = 0$. In contrast, when (1) does not hold, the principal makes a low offer that both types reject.

Suppose next that $b_0 = b_H$. In this case, the assumption that $\mu < \overline{\mu}$ implies that it is optimal for the principal to make a pooling offer $T_0 = c_H$ that both types accept. In particular, if the benefit is large with probability 1 (i.e., $q = 0$), the principal is never able to learn the agent's type.

There are three main takeaways from this example. First, productivity shocks may enable the principal to learn the agent's private information. Second, learning happens when times are bad and the future looks dim. Third, there is path dependence: the value that the principal derives in the second period depends on the first period shock.

In the rest of the paper, we consider an infinite horizon model in which both the agent's type and the principal's benefit can take finitely many values. The three main takeaways of the two-period model extend to this environment. But the infinite horizon model gives rise to new results as well. First, the principal may learn the agent's private information gradually over time. Second, even when learning takes place, learning may stop before the principal knows the agent's type. And finally, the principal's payoff may display path dependence even in the long run and even when the process governing the evolution of productivity is ergodic.

3. MODEL

3.1. *Setup*

We study a repeated game played between a principal and an agent. Time is discrete and indexed by $t = 0, 1, 2, \dots, \infty$. At the start of each period t, a state b_t is drawn from a finite set of states $\mathcal{B}$ and is publicly revealed. The evolution of b_t is governed by a Markov process with transition matrix $[Q_{b,b'}]_{b,b'\in\mathcal{B}}$. After observing $b_t \in \mathcal{B}$, the principal decides how much transfer $T_t \geq 0$ to offer the agent in exchange for taking a productive action. The agent then decides whether or not to take the action. We denote the agent's choice by $a_t \in \{0, 1\}$, where $a_t = 1$ means that the agent takes the action at period t. The action provides the principal a benefit equal to b_t.

The agent incurs a cost $ac \geq 0$ when choosing action $a \in \{0, 1\}$. The agent's cost c of taking the action is his private information and it is fixed throughout the game. Cost c may take one of K possible values from the set $\mathcal{C} = \{c_1, \dots, c_K\}$. The principal's prior belief about the agent's cost is denoted $\mu_0 \in \Delta(\mathcal{C})$, which we assume has full support. At the end of each period the principal observes the agent's action and updates her beliefs about the agent's cost. Players receive their payoffs and the game moves to the next period.[4] Both players are risk-neutral expected utility maximizers and share a common discount factor $\delta < 1$.[5] The payoffs to the principal and an agent of cost $c = c_k$ at the end of period t are, respectively,

$$u(b_t, T_t, a_t) = (1 - \delta)(b_t - T_t)a_t,$$
$$v_k(b_t, T_t, a_t) = (1 - \delta)(T_t - c_k)a_t.$$

[4]As in Hart and Tirole (1988) and Schmidt (1993), the principal can commit to paying the transfer within the current period, but cannot commit to a schedule of transfers in future periods.

[5]The results are qualitatively the same when the players have different discount factors.

We assume, without loss of generality, that the agent's possible costs are ordered so that $0 < c_1 < c_2 < \cdots < c_K$. To avoid having to deal with knife-edge cases, we further assume that $b \neq c_k$ for all $b \in \mathcal{B}$ and $c_k \in \mathcal{C}$. Then it is socially optimal for an agent with cost c_k to take action $a = 1$ at state $b \in \mathcal{B}$ if and only if $b - c_k > 0$. Let the set of states at which it is socially optimal for an agent with cost c_k to take the action be

$$E_k := \{b \in \mathcal{B} : b > c_k\}.$$

We refer to E_k as the *efficiency set* for type c_k. Note that by our assumptions on the ordering of types, the efficiency sets are nested, that is, $E_{k'} \subseteq E_k$ for all $k' \geq k$.

We assume that process $\{b_t\}$ is persistent and that players are moderately patient. To formalize this, first define the following function: for any $b \in \mathcal{B}$ and $B \subseteq \mathcal{B}$, let

$$X(b, B) := (1 - \delta)\mathbb{E}\left[\sum_{t=1}^{\infty} \delta^t \mathbf{1}_{\{b_t \in B\}} \middle| b_0 = b\right],$$

where $\mathbb{E}[\cdot|b_0 = b]$ denotes the expectation operator with respect to the Markov process $\{b_t\}$, given that the period 0 state is b. Thus $X(b, B)$ is the expected discounted amount of time that the realized state is in B in the future, given that the current state is b. For any $b \in \mathcal{B}$, let $b^+ := \{b' \in \mathcal{B} : b' \geq b\}$. We maintain the following assumption throughout.

ASSUMPTION 1—Discounting/Persistence: *For all* $b \in \mathcal{B}$, $X(b, b^+) > 1 - \delta$.

When there are no shocks to productivity (i.e., when the state is fully persistent), this assumption holds when $\delta > 1/2$. In general, for any $\delta > 1/2$, it holds whenever the process $\{b_t\}$ is sufficiently persistent. When process $\{b_t\}$ is ergodic, there is a cutoff $\overline{\delta} \in (1/2, 1)$ such that the assumption holds whenever $\delta > \overline{\delta}$.[6]

3.2. *Histories, Strategies, and Equilibrium Concept*

A history $h_t = \langle(b_0, T_0, a_0), \ldots, (b_{t-1}, T_{t-1}, a_{t-1})\rangle$ records the states, transfers, and agent's action from the beginning of the game until the start of period t. For any two histories $h_{t'}$ and h_t with $t' \geq t$, we write $h_{t'} \succeq h_t$ if the first t period entries of $h_{t'}$ are the same as the t period entries of h_t. Let H_t denote the set of histories of length t and let $H = \bigcup_{t \geq 0} H_t$ denote the set of all histories. A pure strategy for the principal is a function $\tau : H \times \mathcal{B} \to \mathbb{R}_+$, which maps histories and the current state to transfer offers T. A pure strategy for the agent is a collection of mappings $\{\alpha_k\}_{k=1}^K$, $\alpha_k : H \times \mathcal{B} \times \mathbb{R}_+ \to \{0, 1\}$, each of which maps the current history, current state, and current transfer offer to the action choice $a \in \{0, 1\}$ for a particular type c_k.

For conciseness, we restrict attention to pure strategy perfect Bayesian equilibrium (PBE) in the body of the paper. We consider mixed strategies in the Supplemental Material (Acharya and Ortner (2017, Appendix SA.2)); see also Remark 2 below. Pure strategy PBE are denoted by the pair (σ, μ), where $\sigma = (\tau, \{\alpha_k\}_{k=1}^K)$ is a strategy profile and $\mu : H \to \Delta(\mathcal{C})$ gives the principal's beliefs about the agent's type after each history. For any PBE (σ, μ), the continuation payoffs of the principal and an agent with cost c_k after

[6]When there are no shocks, Assumption 1 (i.e., $\delta > 1/2$) guarantees that there will be no learning by the principal in equilibrium; see Hart and Tirole (1988) and Schmidt (1993). If this assumption fails, then there will be learning.

history $h \in H$ and shock realization $b \in \mathcal{B}$ are denoted $U^{(\sigma,\mu)}[h, b]$ and $V_k^{(\sigma,\mu)}[h, b]$. For any $\mu_0 \in \Delta(\mathcal{C})$, any PBE (σ, μ), and any shock $b \in \mathcal{B}$, we denote by $W^{(\sigma,\mu)}[\mu_0, b]$ the principal's payoff at the start of a game with prior μ_0 under the PBE (σ, μ) when the initial shock is b.

We restrict attention to pure strategy PBE that satisfy a sequential optimality condition for the principal, defined as follows. For each integer $n \le |\mathcal{C}| = K$, define $S_n := \{\lambda \in \Delta(\mathcal{C}) : |\operatorname{supp} \lambda| = n\}$. Let Σ_0 denote the set of pure strategy PBE. For all $k = 1, 2, \dots, K$, we define the sets Σ_k recursively as

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

Thus, Σ_1 is the set of pure strategy PBEs that deliver the highest possible payoff to the principal at histories at which her beliefs are degenerate. For all $k > 1$, Σ_k is the set of pure strategy PBEs in Σ_{k-1} that deliver the highest possible payoff to the principal (among all PBEs in Σ_{k-1}) at histories at which the support of her beliefs contains k elements.[7] In what follows, we restrict attention to PBE in Σ_K (recall that $|\mathcal{C}| = K$) and use the word *equilibrium* to refer to such a PBE.[8]

At any PBE satisfying the refinement, the principal extracts the maximum possible surplus from the agent at histories at which she has learned the agent's type; that is, the continuation PBE is "principal-optimal" at such histories. At histories at which the principal believes that the agent may be one of two possible types, players play a continuation PBE that is optimal for the principal among the set of PBEs that are principal-optimal at histories after which the principal has learned the agent's type. We proceed like this to construct restrictions on PBE for continuation games in which the support of the principal's beliefs contain three types, etcetera. By forcing continuation play to be constrained-optimal for the principal at all histories (either on or off the path of play), this solution concept naturally captures lack of commitment by the principal.[9]

We end this section by noting that our equilibrium refinement facilitates a direct comparison with prior papers on the ratchet effect (e.g. Hart and Tirole (1988) and Schmidt (1993)). As we show below, this refinement selects a unique equilibrium that naturally generalizes the equilibrium studied in these papers. In particular, when there are no productivity shocks (i.e., when $\mathcal{B}$ is a singleton), our equilibrium coincides with the equilibrium in Hart and Tirole (1988) and Schmidt (1993).

4. EQUILIBRIUM ANALYSIS

4.1. *Incentive Constraints*

Fix an equilibrium $(\sigma, \mu) = ((\tau, \{\alpha_k\}_{k=1}^K), \mu)$. Recall that for any $h \in H$, $\mu[h]$ are the principal's beliefs at history h. We use $C[h] \subset \mathcal{C}$ to denote the support of $\mu[h]$, and use

[7]We highlight that, for a given belief $\mu[h]$ and state $b \in \mathcal{B}$, we consider only $(\sigma', \mu') \in \Sigma_{k-1}$ such that (σ', μ') is a PBE of a game in which the principal's prior belief is $\mu[h]$ and the initial state is b.

[8]Our solution concept is similar in spirit to the "ratchet equilibrium" concept used by Gerardi and Maestri (2015).

[9]The solution concept can also be interpreted as capturing renegotiation-proofness in a setting where the principal has all of the bargaining power. Consider a PBE in Σ_K. At histories at which the support of the principal's beliefs contains k elements, she has no incentive to renegotiate the equilibrium to another equilibrium in Σ_{k-1}, since the players will play a (constrained) principal-optimal continuation equilibrium.

$\overline{k}[h] := \max\{k : c_k \in C[h]\}$ to denote the highest type index in $C[h]$. Since $c_1 < \cdots < c_K$, $c_{\overline{k}[h]}$ is the highest cost in the support of $\mu[h]$. Finally, for all $c_k \in \mathcal{C}$, we let $a_{t,k}$ be the random variable indicating the action in $\{0, 1\}$ that an agent of type c_k takes in period t under equilibrium (σ, μ).

For any history h, any pair $c_i, c_j \in C[h]$, and any productivity level $b \in \mathcal{B}$, let

$$V_{i\to j}^{(\sigma,\mu)}[h, b] := (1-\delta)\mathbb{E}^{(\sigma,\mu)}\left[\sum_{t'=t}^{\infty} \delta^{t'-t} a_{t',j}(T_{t'} - c_i)\middle| h_t = h, b_t = b\right]$$

be the expected discounted payoff that an agent with cost c_i obtains at time t after history $h_t = h$ when $b_t = b$ from following the equilibrium strategy of an agent with cost c_j. Here $\mathbb{E}^{(\sigma,\mu)}[\cdot|h, b]$ denotes the expectation over future play under equilibrium (σ, μ) conditional on history h and current shock b. Note that for any $c_i \in C[h]$, the continuation value of an agent with cost c_i at history h and current shock b is simply $V_i^{(\sigma,\mu)}[h, b] := V_{i\to i}^{(\sigma,\mu)}[h, b]$. Also note that

$$\begin{aligned} V_{i\to j}^{(\sigma,\mu)}[h, b] &= (1-\delta)\mathbb{E}^{(\sigma,\mu)}\left[\sum_{t'=t}^{\infty} \delta^{t'-t}\big(a_{t',j}(T_{t'} - c_j) + a_{t',j}(c_j - c_i)\big)\middle| h_t = h, b_t = b\right] \\ &= V_j^{(\sigma,\mu)}[h, b] + (c_j - c_i)A_j^{\sigma}[h, b], \end{aligned} \tag{2}$$

where $V_j^{(\sigma,\mu)}[h, b]$ is type c_j's continuation value at (h, b) and

$$A_j^{(\sigma,\mu)}[h, b] := (1-\delta)\mathbb{E}^{(\sigma,\mu)}\left[\sum_{t'=t}^{\infty} \delta^{t'-t} a_{t',j}\middle| h_t = h, b_t = b\right]$$

is the expected discounted number of times that type c_j takes the productive action after history (h, b) under equilibrium (σ, μ). Equation (2) says that type c_i's payoff from deviating to c_j's strategy can be decomposed into two parts: type c_j's continuation value and an *informational rent* $(c_j - c_i)A_j^{(\sigma,\mu)}[h, b]$, which depends on how frequently c_j is expected to take the action in the future. In any equilibrium (μ, σ),

$$V_i^{(\sigma,\mu)}[h, b] \geq V_{i\to j}^{(\sigma,\mu)}[h, b] \quad \forall (h, b), \forall c_i, c_j \in C[h], \tag{3}$$

which represents the set of incentive constraints that must be satisfied. We then have the following lemma, which we prove in the Supplemental Material. Part (i) says that, in any equilibrium, the highest cost type in the support of the principal's beliefs obtains a continuation payoff equal to zero. Part (ii) says that "local" incentive constraints bind.

LEMMA 0: *Fix an equilibrium (σ, μ) and a history h, and if necessary renumber the types so that $C[h] = \{c_1, c_2, \dots, c_{\overline{k}[h]}\}$ with $c_1 < c_2 < \cdots < c_{\overline{k}[h]}$. Then, for all $b \in \mathcal{B}$, the following statements hold*:

(i) *Type $c_{\overline{k}[h]}$ earns zero payoff*: $V_{\overline{k}[h]}^{(\sigma,\mu)}[h, b] = 0$.

(ii) *If $|C[h]| \geq 2$, $V_i^{(\sigma,\mu)}[h, b] = V_{i\to i+1}^{(\sigma,\mu)}[h, b]$ for all $c_i \in C[h] \setminus \{c_{\overline{k}[h]}\}$.*

The proof of Lemma 0 is in the Appendix SA.1. The result follows from our solution concept: in any PBE satisfying our restrictions, the principal will extract all surplus from an agent that has the highest possible cost. Similarly, the principal will extract all possible surplus (subject to incentive compatibility (IC) constraints) from agents with lower costs.

4.2. *Equilibrium Characterization*

We now describe the (essentially) unique equilibrium in Σ_K. Recall that $c_{\overline{k}[h]}$ is the highest cost in the support of the principal's beliefs at history h and that E_k is the set of productivity levels at which it is socially optimal for type $c_k \in \mathcal{C}$ to take the action.

THEOREM 1: *The set of equilibria is nonempty. In any equilibrium* (σ, μ), *for every history* $(h, b) \in H \times \mathcal{B}$, *the following statements hold*:

(i) *If* $b \in E_{\overline{k}[h]}$, *the principal offers transfer* $T = c_{\overline{k}[h]}$ *and all types in* $C[h]$ *take action* $a = 1$.

(ii) *If* $b \notin E_{\overline{k}[h_t]}$, *there is a threshold type* $c_{k^*} \in C[h]$ *such that types in* $C^- := \{c_k \in C[h] : c_k < c_{k^*}\}$ *take action* $a = 1$, *while types in* $C^+ := \{c_k \in C[h] : c_k \geq c_{k^*}\}$ *take action* $a = 0$. *If* C^- *is nonempty*, *the transfer that the principal offers* (*and that is accepted by types in* C^-) *satisfies*

$$T = c_{j^*} + \frac{1}{1-\delta} V^{(\sigma,\mu)}_{j^* \to k^*}[h, b], \qquad (*)$$

where $c_{j^*} = \max C^-$. *If* $X(b, E_{\overline{k}[h]}) > 1 - \delta$, *set* C^- *is empty.*

Theorem 1 says that at histories (h, b) such that either $b \in E_{\overline{k}[h]}$ or $b \notin E_{\overline{k}[h]}$ and $X(b, E_{\overline{k}[h]}) > 1 - \delta$, all the agent types in $C[h]$ take the same action. Hence, the principal learns nothing about the agent's type at such histories. To understand why, note that at such histories (h, b) players expect the state to be in $E_{\overline{k}[h]}$ frequently in the future: formally, $X(b, E_{\overline{k}[h]}) > 1 - \delta$.[10] Therefore, an agent with cost $c_i < c_{\overline{k}[h]}$ gets large rents by mimicking an agent with cost $c_{\overline{k}[h]}$. Since low-cost types anticipate that the principal will leave them with no future rents if they reveal their private information, the principal is unable to learn.

Equilibrium behavior is, however, quite different at histories (h, b) with $b \in E_{\overline{k}[h_t]}$ compared to histories (h, b) with $b \notin E_{\overline{k}[h]}$ and $X(b, E_{\overline{k}[h]}) > 1 - \delta$. When $b \in E_{\overline{k}[h]}$, there is an *efficient* ratchet effect. At these productivity levels the agent takes the socially efficient action $a = 1$ and the principal compensates him as if he was the highest-cost type. This replicates the main finding of the ratchet effect literature. For example, Hart and Tirole (1988) and Schmidt (1993) consider a special case of our model in which the benefit from taking the action is constant over time and strictly larger than the highest cost (i.e., for all times t, $b_t = b > c_K$). Thus, part (i) of Theorem 1 applies: the principal offers a transfer $T = c_K$ that all agent types accept at all periods and she never learns anything about the agent's type.[11]

At histories (h, b) with $b \notin E_{\overline{k}[h]}$ and $X(b, E_{\overline{k}[h]}) > 1 - \delta$, there is an *inefficient* ratchet effect. At these histories, low-cost types pool with high-cost types and do not take the productive action even if the principal is willing to fully compensate their costs. This contrasts with the results in Hart and Tirole (1988) and Schmidt (1993), where the equilibrium is always socially optimal.

At histories (h, b) with $b \notin E_{\overline{k}[h]}$ and $X(b, E_{\overline{k}[h]}) \leq 1 - \delta$, learning may take place. Specifically, the principal learns about the agent's type when a subset of the types take the action (i.e., when the set C^- is nonempty). Intuitively, the informational rent that

[10]For histories (h, b) with $b \in E_{\overline{k}[h_t]}$, this inequality follows from Assumption 1.

[11]Hart and Tirole (1988) and Schmidt (1993) consider games with a finite deadline. In such games, the principal is only able to induce information revelation at the very last periods prior to the deadline. As the deadline grows to infinity, there is no learning by the principal along the equilibrium path.

type $c_i < c_{\overline{k}[h]}$ gets from mimicking an agent with the highest cost $c_{\overline{k}[h]}$ are small when $X(b, E_{\overline{k}[h]}) \leq 1 - \delta$. As a result, the principal is able to get low-cost types to reveal their private information. In Appendix A.1.3, we provide a characterization of the threshold cost c_{k^*} in part (ii) of the theorem as the solution to a finite maximization problem. Building on this, we also characterize the principal's equilibrium payoffs as the fixed point of a contraction mapping.

REMARK 1—Markovian Equilibrium: Note that the equilibrium characterized in Theorem 1 is Markovian: at each period, the behavior of principal and agent depends solely on the principal's beliefs and the current shock realization.

REMARK 2—Mixed Strategies: In the Appendix SA.2, we extend our analysis and consider a broad class of mixed strategy PBE. In particular, we look at the class of finitely revealing PBEs (Peski (2008)); that is, PBEs in which, along any history, the principal's beliefs are updated only finitely many times.

Let Σ_0^M denote the set of PBE that are finitely revealing. For $k = 1, \dots, K$, define the sets Σ_k^M recursively as

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

This is the natural generalization of our equilibrium refinement to mixed strategies under which the principal updates her beliefs a bounded number of times at every history.

Let (σ^P, μ^P) denote the PBE in Theorem 1. We show in Appendix SA.2 that $(\sigma^P, \mu^P) \in \Sigma_K^M$. This implies that any equilibrium in Σ_K^M must give the principal the same payoff as (σ^P, μ^P) at every history. Moreover, we show along the way that generically any equilibrium in the set Σ_K^M is outcome-equivalent to (σ^P, μ^P).

REMARK 3—Full-Commitment Benchmark: We can compare the principal's equilibrium profits to what she would obtain if she had full commitment. A principal with commitment power will in general want to make a high-cost agent take action $a = 1$ inefficiently few times to reduce the informational rents of low-cost types. Time-varying shocks enable the principal to approximate the full-commitment solution. At histories (h, b) with $b \notin E_{\overline{k}[h]}$ and $X(b, E_{\overline{k}[h]}) \leq 1 - \delta$, the principal can truthfully commit to contract infrequently with the highest-cost agent $c_{\overline{k}[h]}$ in the future. This reduces the rents for lower-cost types and enables the principal to learn about the agent's type.

In Appendix SA.3, we illustrate this for the case of two types, $\mathcal{C} = \{c_1, c_2\}$. We show that if $X(b, E_2) = \varepsilon \leq 1 - \delta$ for some productivity level $b \in E_1 \setminus E_2$, then the principal's equilibrium payoff at histories (h, b) with $C[h] = \mathcal{C}$ are within $\varepsilon(1 - \mu)(c_2 - c_1)$ of her full commitment payoff, where $\mu \in (0, 1)$ is the prior probability that $c = c_2$.

4.3. *Examples*

We end this section with two examples that illustrate some of the main equilibrium features of our model. The first highlights the fact that the equilibrium outcome can be inefficient. The second illustrates a situation in which the principal learns the agent's type and the equilibrium outcome is efficient.

EXAMPLE 1—Inefficient Ratchet Effect: Suppose that there are two states, $\mathcal{B} = \{b_L, b_H\}$, and two types, $\mathcal{C} = \{c_1, c_2\}$ with $c_1 < b_L < c_2 < b_H$, so that $E_1 = \{b_L, b_H\}$ and $E_2 = \{b_H\}$. Assume further that $X(b_L, \{b_H\}) > 1 - \delta$.

Consider a history h_t such that $C[h_t] = \{c_1, c_2\}$. Theorem 1(i) implies that, at such a history, both types take the action if $b_t = b_H$, receiving a transfer equal to c_2. On the other hand, Theorem 1(ii) implies that neither type takes the action if $b_t = b_L$. Indeed, when $X(b_L, \{b_H\}) > 1 - \delta$, the benefit that a c_1 agent obtains by pooling with a c_2 agent is so large that there does not exist an offer that a c_1 agent would accept but a c_2 agent would reject. As a result, the principal never learns the agent's type. Inefficiencies arise in all periods t in which $b_t = b_L$: an agent with cost c_1 never takes the action when the state is b_L, even though it is socially optimal for him to do so.

EXAMPLE 2—Efficiency and Learning: The environment is the same as in Example 2, with the only difference that $X(b_L, \{b_H\}) < 1 - \delta$. Consider a history h_t such that $C[h_t] = \{c_1, c_2\}$. As in Example 1, both types take the action in period t if $b_t = b_H$. The difference is that if $b_t = b_L$, the principal offers a transfer T_t that a c_2 agent rejects, but a c_1 agent accepts. The principal's offer T_t exactly compensates type c_1 for revealing his type: $(1 - \delta)(T_t - c_1) = X(b_L, \{b_H\})(c_2 - c_1)$.[12] Note that $X(b_L, \{b_H\}) < 1 - \delta$ implies that $T_t < c_2$, so an agent with cost c_2 rejects offer T_t. The principal finds it optimal to make such an offer, since it gets an agent with cost $c_L < b_L$ to take the efficient action.

We note that the principal learns the agent's type at time $\bar{t} = \min\{t : b_t = b_L\}$, and the outcome is efficient from time $\bar{t} + 1$ onward: type c_i takes the action at time $t' > \bar{t}$ if and only if $b_{t'} \in E_i$. Moreover, Lemma 0(i) guarantees that the principal extracts all of the surplus from time $\bar{t} + 1$ onward, paying the agent a transfer equal to his cost.

The inefficiency in Example 1 contrasts with the results of the ratchet effect literature in which the outcome is always efficient. The results of Example 2 contrast with this literature as well, in which the principal finds it difficult to learn. The key features of this example are that (i) learning by the principal takes place only if productivity is low, (ii) the principal eventually achieves her first best payoff, and (iii) the equilibrium exhibits a form of path dependence: equilibrium play at time t depends on the entire history of shocks up to period t.[13] These features motivate the results of the next section.

5. IMPLICATIONS

5.1. *The Consequences of Bad Shocks*

In Example 2 above, the principal learns the agent's type and learning takes place the first time the relationship hits the low productivity state. In addition, as soon as the low productivity state is reached for the first time, the agent's compensation falls permanently. In this section, we present conditions under which these results generalize.

Consider the following assumption, which is a monotonicity condition on the stochastic process that governs the evolution of productivity.

[12] The payoff a low-cost agent obtains by accepting offer T_t is $(1 - \delta)(T_t - c_1) + \delta 0$, since the principal learns that the agent's cost is c_1. On the other hand, the payoff such an agent obtains from rejecting the offer and mimicking a high-cost agent is $X(b_L, \{b_H\})(c_2 - c_1)$.

[13] Before state b_L is reached for the first time, the principal pays a transfer equal to the agent's highest cost c_2 to get both types to take the action. After state b_L is visited, if the principal finds that the agent has low cost, then she pays a lower transfer equal to c_1.

ASSUMPTION 2: *For all* $c_k \in \mathcal{C}$, $X(b, E_k) \leq X(b', E_k)$ *for all* $b, b' \in \mathcal{B}$ *with* $b < b'$.

The assumption is natural; for example, it holds when transition matrix $\{Q_{b,\tilde{b}}\}_{b,\tilde{b}\in\mathcal{B}}$ satisfies the monotone likelihood ratio property.[14]

Now refer to history (h_t, b_t) as a *history of information revelation* if $\mu[h_{t+1}] \neq \mu[h_t]$; that is, if learning takes place at history (h_t, b_t). The following proposition states that under Assumption 2, learning takes place only in periods of low productivity.

PROPOSITION 1—Learning in Bad Times: *Suppose that Assumption 2 holds. For every history* h_t *there exists a productivity level* $b[h_t] \in \mathcal{B}$ *such that* (h_t, b_t) *is a history of information revelation only if* $b_t < b[h_t]$.

PROOF: By Theorem 1, $\mu[h_{t+1}] \neq \mu[h_t]$ only if (h_t, b_t) are such that $X(b_t, E_{\overline{k}[h_t]}) \leq 1 - \delta$. By Assumption 2, there exists $b[h_t]$ such that $X(b_t, E_{\overline{k}[h_t]}) \leq 1 - \delta$ if and only if $b_t < b[h_t]$.[15] *Q.E.D.*

To see why the result holds, note that under Assumption 2 the future expected discounted surplus of the relationship is increasing in the current shock b_t. This implies that the informational rent that agents with type $c_i < c_{\overline{k}[h_t]}$ get from mimicking an agent with the highest cost $c_{\overline{k}[h_t]}$ is also increasing in b_t. As a result, the principal is only able to learn about the agent's type in periods where the productivity b_t is low.

Next, recall that according to Theorem 1, if (h_t, b_t) is a history of information revelation, then there exists a type $c_{j^*} \in C[h_t]$ such that only agents with cost at most c_{j^*} take action at time t. We refer to type c_{j^*} as the *marginal type* in period t. Also, for every history (h_t, b_t) and every type $c_j \in C[h_t]$, define

$$P_j[h_t, b_t] := -(1-\delta)\mathbb{E}^{(\sigma,\mu)}\left[\sum_{t'=t}^{\infty} \delta^{t'-t} |\mathbf{1}_{b_{t'}\in E_j} - a_{t',j}|(b_{t'} - c_j)\middle| h_t, b_t\right],$$

which is a measure of how efficient the equilibrium actions of type c_j are. The following proposition, which follows directly from Theorem 1, states two results: (i) that productivity increases after histories of information revelation and (ii) that the agent's compensation may fall permanently after such histories.

PROPOSITION 2—Productivity and Compensation: *Let* (h_t, b_t) *be a history of information revelation, and let* c_{j^*} *be the marginal type at time* t. *Then, for all* (h_τ, b_τ) *with* $h_\tau \succeq h_t$,
(i) $P_{j^*}[h_\tau, b_\tau] = 0$,
(ii) $V_{j^*}^{(\sigma,\mu)}[h_\tau, b_\tau] = 0$.

Part (i) of this result, combined with Proposition 1, implies that agents' productivity will increase after the relationship goes through bad times. The result is in line with Lazear, Shaw, and Stanton (2016), who find evidence that workers' productivity increases after a recession. Part (ii) combined with Proposition 1 implies that the agents' compensation

[14]That is, for every $b > b'$, $\frac{Q_{b,\tilde{b}}}{Q_{b',\tilde{b}}} = \frac{\text{prob}(b_{t+1}=\tilde{b}|b_t=b)}{\text{prob}(b_{t+1}=\tilde{b}|b_t=b')}$ is increasing in $\tilde{b}$.

[15]When $b[h_t] = \min \mathcal{B}$, $X(b, E_k) > 1 - \delta$ for all $b \in \mathcal{B}$. In this case, the principal's beliefs remain unchanged after history h_t.

may be permanently lowered after the relationship experiences negative shocks. This finding is consistent with Kahn (2010) and Oreopoulos, Von Wachter, and Heisz (2012), who provide evidence that recessions have a persistent negative effect on worker compensation.

We end this section by briefly discussing the robustness of these predictions to settings in which both the firm's productivity and the worker's outside option are time-varying and publicly observed. In such a setting, learning will typically take place at periods in which net productivity (i.e., productivity minus outside option) is low. Therefore, our predictions relating recessions to the agent's performance and compensation would continue to hold as long as recessions negatively effect the firm's net productivity.

5.2. *Long-Run First-Best Payoffs*

Another notable feature of Example 2 is that full learning takes place and, as a result, the principal's value increases permanently to the first-best level. Here, we characterize general conditions under which the principal obtains her first-best payoff in the long run, as well as conditions under which she does not. Before stating our results, we introduce some additional notation and make a preliminary observation.

An equilibrium outcome can be written as an infinite sequence $h_\infty = \langle b_t, T_t, a_t\rangle_{t=0}^\infty$ or, equivalently, as an infinite sequence of equilibrium histories $h_\infty = \{h_t\}_{t=0}^\infty$ such that $h_{t+1} \succeq h_t$ for all t. For any equilibrium outcome h_∞, there exists a time $t^*[h_\infty]$ such that $\mu[h_t] = \mu[h_{t^*[h_\infty]}]$ for all $h_t \succeq h_{t^*[h_\infty]}$. That is, given an equilibrium outcome, learning always stops after some time $t^*[h_\infty]$. Given an equilibrium outcome h_∞, in every period after $t^*[h_\infty]$, the principal's continuation payoff depends only on the realization of the current period shock. Formally, given any equilibrium outcome $h_\infty = \{h_t\}_{t=0}^\infty$, the principal's equilibrium continuation value at time $t \geq t^*[h_\infty]$ can be written as $U_{\text{LR}}^{(\sigma,\mu)}(b_t|h_{t^*[h_\infty]})$.

For all $b \in \mathcal{B}$ and all $c_k \in \mathcal{C}$, the principal's first-best payoffs conditional on the current shock being b and the agent's type being $c = c_k$ are given by

$$U^*(b|c_k) := (1-\delta)\mathbb{E}\left[\sum_{t'=t}^{\infty} \delta^{t'-t}(b_{t'} - c_k)\mathbf{1}_{\{b_{t'} \in E_k\}} \middle| b_t = b\right].$$

Under the first-best outcome, the agent takes the action whenever it is socially optimal and the principal always compensates the agent his exact cost. Say that an equilibrium (σ, μ) is *long-run first-best* if for all $c_k \in \mathcal{C}$, the set of equilibrium outcomes h_∞ such that

$$U_{\text{LR}}^{(\sigma,\mu)}(b|h_{t^*[h_\infty]}) = U^*(b|c_k) \quad \forall b \in \mathcal{B}$$

has probability 1 when the agent's type is $c = c_k$. The next result, which we prove in Appendix A.2, reports a sufficient condition for the equilibrium to be long-run first-best.

PROPOSITION 3—Long-Run First-Best: *Suppose that $\{b_t\}$ is ergodic and that for all $c_k \in \mathcal{C} \setminus \{c_K\}$ there exists a productivity level $b \in E_k \setminus E_{k+1}$ such that $X(b, E_{k+1}) < 1 - \delta$. Then the equilibrium is long-run first-best.*

The condition in the statement of Proposition 3 guarantees that, for any history h such that $|C[h]| \geq 2$, there exists at least one state $b \in \mathcal{B}$ at which the principal finds it optimal to make an offer that only a strict subset of types accept. So if the process $\{b_t\}$ is ergodic, then it is certain that the principal will eventually learn the agent's type, and from that point onward she gets her first-best payoffs.

If an equilibrium is long-run first-best, then it is also *long-run efficient*, that is, for all $c_k \in \mathcal{C}$, with probability 1 an agent with cost c_k takes the action at time $t > t^*[h_\infty]$ if and only if $b_t \in E_k$. However, the converse of this statement is not true. Because of this, there are weaker sufficient conditions under which long-run efficiency holds. One such condition is that $\{b_t\}$ is ergodic and for all $c_k \in \mathcal{C}$ such that $E_k \neq E_K$, there exists $b \in E_k \setminus E_{\underline{k}}$ such that $X(b, E_{\underline{k}}) < 1 - \delta$, where $\underline{k} = \min\{j \geq k : E_j \neq E_k\}$. This condition guarantees that the principal's beliefs will eventually place all the mass on the set of types that share the same efficiency set with the agent's true type. After this happens, even if the principal does not achieve her first-best payoff by further learning the agent's type, the agent takes the action if and only if it is socially optimal to do so. The argument mirrors that of Proposition 3.

Our next result provides a partial counterpart to Proposition 3. The result is an immediate consequence of Theorem 1.

PROPOSITION 4—No Long-Run First-Best; No Long-Run Efficiency: *Let h be an equilibrium history such that $|C[h]| \geq 2$ and suppose that $X(b, E_{\overline{k}[h]}) > 1 - \delta$ for all $b \in \mathcal{B}$. Then $\mu[h'] = \mu[h]$ for all histories $h' \succeq h$ (and thus $|C[h']| \geq 2$), so the equilibrium is not long-run first-best. If, in addition, there exists $c_i \in C[h]$ such that $E_i \neq E_{\overline{k}[h]}$, then the equilibrium is not long-run efficient either.*

5.3. *Long-Run Path Dependence*

The third notable feature of Example 2 was that the equilibrium exhibits a form of path dependence: equilibrium play at time t depends on the entire history of shocks up to period t. Note, however, that the path dependence in Example 2 is short-lived: after state b_L is visited for the first time, the principal learns the agent's type and the equilibrium outcome from that point onward is independent of the prior history of shocks. Here we show that this is *not* a general property of our model.

Say that an equilibrium (σ, μ) exhibits *long-run path dependence* if for some type of the agent $c_k \in \mathcal{C}$ there exists $U_1 : \mathcal{B} \to \mathbb{R}$ and $U_2 : \mathcal{B} \to \mathbb{R}$, $U_1 \neq U_2$, such that conditional on the agent's type being c_k, the set of outcomes h_∞ with $U_{\text{LR}}^{(\sigma,\mu)}(\cdot|h_{t^*[h_\infty]}) = U_i(\cdot)$ has positive probability for $i = 1, 2$. That is, the equilibrium exhibits long-run path dependence if, with positive probability, the principal's long-run payoffs may take more than one value conditional on the agent's type.

The next example shows that equilibrium may exhibit long-run path dependence when process $\{b_t\}$ is not ergodic.

EXAMPLE 3—Path Dependence With Non-Ergodic Shocks: Let $\mathcal{C} = \{c_1, c_2\}$, and $\mathcal{B} = \{b_L, b_M, b_H\}$, with $b_L < b_M < b_H$. Suppose that $E_1 = \{b_L, b_M, b_H\}$ and $E_2 = \{b_M, b_H\}$. Suppose further that the transition matrix $[Q_{b,b'}]$ satisfies (i) $X(b_L, E_2) < 1 - \delta$ and (ii) $Q_{b_H,b_H} = 1$ and $Q_{b,b'} \in (0, 1)$ for all (b, b') with $b \neq b_H$. Thus, state b_H is absorbing. By Theorem 1, if $b_t = b_H$, from time t onward the principal makes an offer equal to $c_{\overline{k}[h_t]}$ and all agent types in $C[h_t]$ accept.

Consider history h_t with $C[h_t] = \{c_1, c_2\}$. By Theorem 1, if $b_t = b_M$ the principal makes an offer $T_t = c_2$ that both types of agents accept. If $b_t = b_L$, the principal makes offer $\tilde{T} = c_1 + \frac{1}{1-\delta} X(b_L, E_2)(c_2 - c_1) \in (c_1, c_2)$ that type c_1 accepts and type c_2 rejects. Therefore, the principal learns the agent's type.

Now suppose that the agent's true type is $c = c_1$, and consider the two histories, h_t and $\tilde{h}_t$:

$$h_t = \langle (b_{t'} = b_M, T_{t'} = c_2, a_{t'} = 1)_{t'=1}^{t-1} \rangle,$$

$$\tilde{h}_t = \langle (b_{t'} = b_M, T_{t'} = c_2, a_{t'} = 1)_{t'=1}^{t-2}, (b_{t-1} = b_L, T_{t-1} = \tilde{T}, a_{t-1} = 1) \rangle.$$

Under history h_t, $b_{t'} = b_M$ for all $t' \leq t - 1$, so the principal's beliefs after h_t is realized are equal to her prior. Under history $\tilde{h}_t$, the principal learns that the agent's type is c_1 at time $t - 1$. Suppose that $b_t = b_H$, so that $b_{t'} = b_H$ for all $t' \geq t$. Under history h_t, the principal does not know the agent's type at t, and therefore offers a transfer $T_{t'} = c_2$ for all $t' \geq t$, which both agent types accept. However, under history $\tilde{h}_t$, the principal knows that the agent's type is c_1 and therefore offers transfer $T_{t'} = c_1$ for all $t' \geq t$, and the agent accepts it. Therefore, when the agent's type is c_1, the principal's continuation payoff at history $(h_t, b_t = b_H)$ is $b_H - c_2$, while her payoff at history $(\tilde{h}_t, b_t = b_H)$ is $b_H - c_1$.

Path dependence in this example is driven by the non-ergodicity of the productivity shocks. Since $b_H > c_2$ is absorbing, Theorem 1 implies that the principal will stop learning once the shock reaches this state. At the same time, the principal is able to screen the different types when the shock reaches state b_L (since $X(b_L, E_2) < 1 - \delta$), but is unable to screen them at state b_M. Therefore, the principal only learns the agent's type at histories such that shock b_L is realized before shock b_H.

We highlight, however, that the model may give rise to path dependence even when the evolution of productivity is governed by an ergodic process. The following example, which is fully developed in Appendix SA.4, illustrates this.

EXAMPLE 4—Path Dependence With Ergodic Shocks: Let $\mathcal{C} = \{c_1, c_2, c_3\}$ and $\mathcal{B} = \{b_L, b_{ML}, b_{MH}, b_H\}$, $b_L < b_{ML} < b_{MH} < b_H$. Suppose that $E_1 = E_2 = \{b_{ML}, b_{MH}, b_H\}$ and $E_3 = \{b_H\}$. Suppose further that the transition matrix $[Q_{b,b'}]$ satisfies (a) $Q_{b,b'} > 0$ for all $b, b' \in \mathcal{B}$, and (b) $X(b_{MH}, \{b_H\}) > 1 - \delta$ and $X(b, \{b_H\}) < 1 - \delta$ for $b = b_L, b_{ML}$.

In Appendix SA.4 we show that, under additional conditions, the unique equilibrium has 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 $a = 1$.

An immediate consequence of these facts is that when the agent's type is c_1, the principal learns the agent's type at histories such that state b_L is visited before b_{ML}. In contrast, at histories at which b_{ML} is visited before b_L, the principal only learns that the agent's type is in $\{c_1, c_2\}$. From this point onward, her beliefs are never again updated. As a result, the principal's long-run value when the agent's type is c_1 depends on whether or not shock b_L is realized before shock b_{ML}.

To understand Example 4, note that the informational rents that type c_1 gets by mimicking type c_2 depend on how often c_2 is expected to take the productive action in the future.

In turn, how often c_2 takes the productive action depends on the principal's beliefs. If the principal learns along the path of play that the agent's type is not c_3, from that time onward type c_2 will take the action whenever the state is in $E_2 = \{b_{ML}, b_{MH}, b_H\}$.

In contrast, at histories at which the principal has not ruled out types c_2 and c_3, type c_2 will not take the productive action at time t if $b_t = b_{MH}$ (since, by assumption, $X(b_{MH}, E_3) > 1 - \delta$). Therefore, type c_2 is expected to take the action significantly less frequently in the future at a history after which the support of the principal's beliefs is $\{c_1, c_2, c_3\}$ than at a history at which it is $\{c_1, c_2\}$.

As a consequence of this, the cost of getting a c_1 agent to reveal his private information depends on the principal's beliefs. In particular, when the current productivity level is b_L, getting a c_1 agent to reveal his private information is cheaper at histories where all three types are in the support of the principal's beliefs than at histories at which only c_1 and c_2 are in the support. This difference makes it optimal for the principal to get a c_1 agent to reveal his type when productivity is b_L and the support of the principal's beliefs is $\{c_1, c_2, c_3\}$, and at the same time it makes it suboptimal to get this agent type to reveal himself when productivity is b_L and the support is $\{c_1, c_2\}$.

6. FINAL REMARKS

Productivity shocks are a natural feature of most economic environments, and the incentives that economic agents face in completely stationary environments can be very different than the incentives they face in environments subject to these shocks. Our results demonstrate this for the traditional ratchet effect literature. A key takeaway from this literature is that outside institutions that provide contract enforcement can help improve the principal's welfare. However, our results show that even without such institutions, a strategic principal can use productivity shocks to her advantage to progressively learn the agent's private information and improve her own welfare.

Our model has several natural extensions. For example, we have assumed that the benefit b that the principal obtains when the agent takes the action is publicly observed. This assumption is natural in settings in which the principal's benefit depends on the cost of some key input (like oil or cement) or when this benefit is linked to the aggregate state of the economy. However, it is also interesting to consider settings in which benefit b is privately observed by the principal.

For concreteness, consider the setting of Examples 1 and 2, in which b can take values $\{b_L, b_H\}$ and the agent's cost can take values $\{c_1, c_2\}$, with $b_H > c_2 > b_L > c_1$ (i.e., $E_1 = \{b_L, b_H\}$ and $E_2 = \{b_H\}$). Assume first that $X(b, E_2) > 1 - \delta$ for $b = b_L, b_H$. In this case, the equilibrium outcome in Theorem 1 remains an equilibrium even when $\{b_t\}$ is privately observed: under this condition a low-cost agent is not willing to disclose his private information regardless of whether or not he observes the shock realization.

In constrast, if $X(b_L, E_2) < 1 - \delta$, the equilibrium outcome in Theorem 1 fails to be an equilibrium when shocks are privately observed by the principal. In this case, when the benefit is b_H the principal would prefer to make an offer as if the benefit were b_L, to induce the low-cost agent to reveal his private information. In this setting, one can construct PBEs under which the principal's transfer offer perfectly reveals her private information at each point (i.e., her transfer T_t reveals the realization of benefit b_t at every period t). In such equilibria, at histories (h, b) with $b = b_H$ and $C[h] = \{c_1, c_2\}$, the principal makes a high offer $T = c_2$ that both types accept with probability 1. At histories (h, b) with $b = b_L$ and $C[h] = \{c_1, c_2\}$, the principal makes a low offer $T \in (c_1, c_2)$ that leaves low-cost agents indifferent between accepting and rejecting. Such an offer is rejected by high-cost agents and is accepted with probability $\alpha \in [0, 1]$ by low-cost agents. The probability of acceptance α is calibrated to provide incentives to the principal to make a high offer $T = c_2$ at histories (h, b) with $b = b_H$ and $C[h] = \{c_1, c_2\}$.[16]

APPENDIX

A.1. *Proof of Theorem 1*

The proof proceeds in three steps. First we analyze the case where $b_t \in E_{\overline{k}[h_t]}$, establishing part (i) of the theorem. Then we analyze the case where $b_t \notin E_{\overline{k}[h_t]}$, establishing part (ii). Finally, we show that equilibrium exists and has unique payoffs. In doing so, we also characterize the threshold type c_{k^*} defined in part (ii).

A.1.1. *Proof of Part (i) (the Case of $b_t \in E_{\overline{k}[h_t]}$)*

We prove part (i) by strong induction on the cardinality of $C[h_t]$. If $C[h_t]$ is a singleton $\{c_k\}$, the result holds: in any PBE in Σ_K, the principal offers the agent a transfer $T_{t'} = c_k$ at all times $t' \geq t$ such that $b_{t'} \in E_k$ and the agent accepts, and she offers some transfer $T_{t'} < c_k$ at all times $t' \geq t$ such that $b_{t'} \notin E_k$, and the agent rejects.

Suppose next that the claim is true for all histories $h_{t'}$ such that $|C[h_{t'}]| \leq n - 1$. Let (h_t, b_t) be a history such that $|C[h_t]| = n$ and $b_t \in E_{\overline{k}[h_t]}$. We show that at such a history (h_t, b_t), the principal makes an offer $T_t = c_{\overline{k}[h_t]}$ that all agent types accept.

Note first that in a PBE in Σ_K, it cannot be that at (h_t, b_t) the principal makes an offer that no type in $C[h_t]$ accepts. To see this, suppose by contradiction that this were possible. Then consider an alternative PBE that is identical to the original PBE, except that at history (h_t, b_t) the principal makes an offer $T = c_{\overline{k}[h_t]}$, and all agent types in $C[h_t]$ accept any offer weakly larger than $T = c_{\overline{k}[h_t]}$. The principal's beliefs after this period are equal to $\mu[h_t]$ regardless of the agent's action. Since $T = c_{\overline{k}[h_t]}$, it is optimal for all agent types to accept this offer. Moreover, it is optimal for the principal to make offer T. Finally, since $b_t \in E_{\overline{k}[h_t]}$, the payoff that the principal gets from this PBE is larger than her payoff under the original PBE. But this cannot be, since the original PBE is in Σ_K. Hence, if $b_t \in E_{\overline{k}[h_t]}$, at least a subset of types in $C[h_t]$ take the action at time t if $b_t \in E_{\overline{k}[h_t]}$.

We now show that in a PBE in Σ_K, it cannot be that at (h_t, b_t) the principal makes an offer T_t that only a strict subset $C \subsetneq C[h_t]$ of types accept. Toward a contradiction, suppose that a strict subset $C \subsetneq C[h_t]$ of types accept T_t, and let $c_j = \max C$. There are two possible cases: (a) $c_j = c_{\overline{k}[h_t]}$ and (b) $c_j < c_{\overline{k}[h_t]}$. Consider case (a). By Lemma 0, the continuation payoff of an agent with cost $c_{\overline{k}[h_t]}$ is zero at all histories. This implies that $T_t = c_{\overline{k}[h_t]}$. Let $c_i = \max C[h_t] \setminus C$ (note that $C[h_t] \setminus C$ is nonempty by assumption). Since c_i rejects the offer today and becomes the highest cost in the support of the principal's beliefs tomorrow, Lemma 0 implies that $V_i^{(\sigma,\mu)}[h_t, b_t] = 0$. But this cannot be, since this agent can guarantee a payoff of at least $(1 - \delta)(T_t - c_i) = (1 - \delta)(c_{\overline{k}[h_t]} - c_i) > 0$ by accepting the offer. Hence, if only a strict subset $C \subsetneq C[h_t]$ of types accept, $c_j = \max C < c_{\overline{k}[h_t]}$.

Consider next case (b). By Lemma 0, the payoff of type c_j from taking the productive action at time t is $(1 - \delta)(T_t - c_j) + 0$. Indeed, at period $t + 1$, c_j will be the highest cost in the support of the principal's beliefs if he takes the action at t. Since an agent with cost

[16]Further details about such equilibria are available from the authors upon request.

c_j can mimic the strategy of type $c_{\overline{k}[h_t]}$, incentive compatibility implies that

$$\begin{aligned}(1-\delta)(T_t - c_j) &\geq V^{(\sigma,\mu)}_{\overline{k}[h_t]}[h_t, b_t] + (c_{\overline{k}[h_t]} - c_j) A^{(\sigma,\mu)}_{\overline{k}[h_t]}[h_t, b_t] \\ &\geq (c_{\overline{k}[h_t]} - c_j) X(b_t, E_{\overline{k}[h_t]}) > (1-\delta)(c_{\overline{k}[h_t]} - c_j).\end{aligned} \tag{4}$$

The first inequality follows from equation (3) in the main text. The second inequality follows from Lemma 0 and the fact that $A^{(\sigma,\mu)}_{\overline{k}[h_t]}[h_t, b_t] \geq X(b_t, E_{\overline{k}[h_t]})$. To see why this last inequality holds, note that $c_{\overline{k}[h_t]} \notin C$, so at most $n-1$ types accept the principal's offer. Thus, the inductive hypothesis implies that if the agent rejects the offer, then at all periods $t' > t$, the principal will get all the remaining types to take the action whenever $b_t \in E_{\overline{k}[h_t]}$, and so $A^{(\sigma,\mu)}_{\overline{k}[h_t]}[h_t, b_t] \geq X(b_t, E_{\overline{k}[h_t]})$. The last inequality in equation (4) follows from the fact $X(b_t, E_{\overline{k}[h_t]}) \geq X(b_t, b_t^+) > 1-\delta$, where the first inequality holds because $b_t \in E_{\overline{k}[h_t]}$ and the second inequality follows by Assumption 1.

On the other hand, because Lemma 0 implies that an agent with type $c_{\overline{k}[h_t]}$ has a continuation value of zero, the transfer T_t that the principal offers must be weakly smaller than $c_{\overline{k}[h_t]}$; otherwise, if $T_t > c_{\overline{k}[h_t]}$, an agent with type $c_{\overline{k}[h_t]}$ could guarantee himself a strictly positive payoff by accepting the offer. But this contradicts (4). Hence, it cannot be that only a strict subset of types in $C[h_t]$ accept the principal's offer at (h_t, b_t).

By the arguments above, all agents in $C[h_t]$ take action $a = 1$ at (h_t, b_t) with $b_t \in E_{\overline{k}[h_t]}$. Since an agent with cost $c_{\overline{k}[h_t]}$ obtains a payoff of zero after every history (Lemma 0), the transfer that the principal offers at time t is $T_t = c_{\overline{k}[h_t]}$. *Q.E.D.*

A.1.2. *Proof of Part (ii) (the Case of $b_t \notin E_{\overline{k}[h_t]}$)*

At histories (h_t, b_t) with $b_t \notin E_{\overline{k}[h_t]}$, the highest cost type in the principal's support $c_{\overline{k}[h_t]}$ does not take the productive action. We prove this in Lemma A.1 below and use the lemma to prove part (ii) of the theorem.

LEMMA A.1: *Fix any equilibrium (σ, μ) and history h_t. If $b_t \notin E_{\overline{k}[h_t]}$, then an agent with cost $c_{\overline{k}[h_t]}$ does not take the productive action at time t.*

PROOF: Suppose for the sake of contradiction that an agent with type $c_{\overline{k}[h_t]}$ does take the action at time t if $b_t \notin E_{\overline{k}[h_t]}$. Since, by Lemma 0, this type's payoff must equal zero at all histories, it must be that the offer that is accepted is $T_t = c_{\overline{k}[h_t]}$. We now show that if the principal makes such an offer, then all agent types will accept the offer and take the productive action. To see this, suppose some types reject the offer. Let c_j be the highest-cost type that rejects the offer. By Lemma 0, type c_j's continuation payoff is zero, because this type becomes the highest cost in the support of the principal's beliefs following a rejection. However, this type can guarantee himself a payoff of at least $(1-\delta)(T_t - c_j) = (1-\delta)(c_{\overline{k}[h_t]} - c_j) > 0$ by accepting the current offer. Hence, it cannot be that some types reject offer $T_t = c_{\overline{k}[h_t]}$ when type $c_{\overline{k}[h_t]}$ accepts it.

It then follows that if type $c_{\overline{k}[h_t]}$ accepts the offer, then the principal will not learn anything about the agent's type. Since $b_t \notin E_{\overline{k}[h_t]}$, her flow payoff from making the offer is $(1-\delta)(b_t - c_{\overline{k}[h_t]}) < 0$. Consider an alternative PBE that is identical to the original PBE except that at history (h_t, b_t), the principal makes an offer $T = 0$ and all agent types in $C[h_t]$ reject this offer. The principal's beliefs after this period are equal to $\mu[h_t]$ regardless of the agent's action. Note that it is optimal for all types to reject this offer. Moreover, since $b_t \notin E_{\overline{k}[h_t]}$, the payoff that the principal gets from this PBE is larger than her payoff under the original PBE. But this cannot be since the original PBE is in Σ_K. Hence, if $b_t \notin E_{\overline{k}[h_t]}$, an agent with type $c_{\overline{k}[h_t]}$ does not take the action at time t. *Q.E.D.*

PROOF OF PART (II) WHEN $X(b_t, E_{\overline{k}[h_t]}) > 1-\delta$: Fix a history (h_t, b_t) with $b_t \notin E_{\overline{k}[h_t]}$ and $X(b_t, E_{\overline{k}[h_t]}) > 1-\delta$. We show that at such histories, all agent types take action $a = 0$.

By Lemma A.1, type $c_{\overline{k}[h_t]}$ does not take the productive action at time t if $b_t \notin E_{\overline{k}[h_t]}$. Suppose, for the sake of contradiction, that there is a nonempty set of types $C \subsetneq C[h_t]$ that do take the productive action. Let $c_j = \max C$. By Lemma 0, type c_j obtains a continuation payoff of zero starting in period $t+1$. Hence, type c_j receives a payoff $(1-\delta)(T_t - c_j) + \delta 0$ from taking the productive action in period t. Since this payoff must be weakly larger than the payoff the agent would obtain by not taking the action and mimicking the strategy of agent $c_{\overline{k}[h_t]}$ in all future periods, it follows that

$$\begin{aligned}(1-\delta)(T_t - c_j) &\geq V^{(\sigma,\mu)}_{\overline{k}[h_t]}[h_t, b_t] + (c_{\overline{k}[h_t]} - c_j) A^{(\sigma,\mu)}_{\overline{k}[h_t]}[h_t, b_t] \\ &\geq (c_{\overline{k}[h_t]} - c_j) X(b_t, E_{\overline{k}[h_t]}) \\ &> (1-\delta)(c_{\overline{k}[h_t]} - c_j),\end{aligned} \tag{5}$$

where the first line follows from incentive compatibility, the second line follows from the fact that $a_{t', \overline{k}[h_t]} = 1$ for all times $t' \geq t$ such that $b_{t'} \in E_{\overline{k}[h_t]}$ (by the result of part (i) proven above), and the third line follows since $X(b_t, E_{\overline{k}[h_t]}) > 1-\delta$ by assumption. The inequalities in (5) imply that $T_t > c_{\overline{k}[h_t]}$. But then by Lemma 0, it would be strictly optimal for type $c_{\overline{k}[h_t]}$ to deviate by accepting the transfer and taking the productive action, a contradiction. So it must be that all agent types in $C[h_t]$ take action $a_t = 0$. *Q.E.D.*

PROOF OF PART (II) WHEN $X(b_t, E_{\overline{k}[h_t]}) \leq 1-\delta$: Fix a history h_t and let $b_t \in \mathcal{B} \setminus E_{\overline{k}[h_t]}$ be such that $X(b_t, E_{\overline{k}[h_t]}) \leq 1-\delta$. We start by showing that the set of types that accept the offer has the form $C^- = \{c_k \in C[h_t] : c_k < c_{k^*}\}$ for some $c_{k^*} \in C[h_t]$. The result is clearly true if no agent type takes the action, in which case set $c_{k^*} = \min C[h_t]$, or if only an agent with type $\min C[h_t]$ takes the action, in which case set c_{k^*} equal to the second lowest cost in $C[h_t]$.

Therefore, suppose that an agent with type larger than $\min C[h_t]$ takes the action, and let $c_{j^*} \in C[h_t]$ be the highest-cost agent that takes the action. Since $b_t \notin E_{\overline{k}[h_t]}$, by Lemma A.1 it must be that $c_{j^*} < c_{\overline{k}[h_t]}$. By Lemma 0, type c_{j^*}'s payoff is $(1-\delta)(T_t - c_{j^*})$, since from date $t+1$ onward this type will be the highest-cost type in the support of the principal's beliefs if the principal observes that the agent took the action at time t. Let $c_{k^*} = \min\{c_k \in C[h_t] : c_k > c_{j^*}\}$. By incentive compatibility, it must be that

$$(1-\delta)(T_t - c_{j^*}) \geq V^{(\sigma,\mu)}_{k^*}[h_t, b_t] + (c_{k^*} - c_{j^*}) A^{(\sigma,\mu)}_{k^*}[h_t, b_t], \tag{6}$$

since type c_{j^*} can obtain the right-hand side of (6) by mimicking type c_{k^*}. Furthermore, type c_{k^*} can guarantee himself a payoff of $(1-\delta)(T_t - c_{k^*})$ by taking the action at time t and never taking the action again. Therefore, it must be that

$$\begin{aligned}V^{(\sigma,\mu)}_{k^*}[h_t, b_t] &\geq (1-\delta)(T_t - c_{k^*}) \\ &\geq (1-\delta)(c_{j^*} - c_{k^*}) + V^{(\sigma,\mu)}_{k^*}[h_t, b_t] + (c_{k^*} - c_{j^*}) A^{(\sigma,\mu)}_{k^*}[h_t, b_t] \\ \implies \quad 1-\delta &\geq A^{(\sigma,\mu)}_{k^*}[h_t, b_t],\end{aligned} \tag{7}$$

where the second inequality in the first line follows from (6).

We now show that all types $c_i \in C[h_t]$ with $c_i < c_{j^*}$ also take the action at time t. Suppose for the sake of contradiction that this is not true, and let $c_{i^*} \in C[h_t]$ be the highest-cost type lower than c_{j^*} that does not take the action. The payoff that this type would get by taking the action at time t and then mimicking type c_{j^*} is

$$\begin{aligned} V_{i^*\to j^*}^{(\sigma,\mu)}[h_t,b_t] &= (1-\delta)(T_t-c_{j^*}) + (c_{j^*}-c_{i^*})A_{j^*}^{(\sigma,\mu)}[h_t,b_t] \\ &= (1-\delta)(T_t-c_{j^*}) + (c_{j^*}-c_{i^*})\big(1-\delta+X(b_t,E_{j^*})\big) \\ &\geq (c_{j^*}-c_{i^*})\big(1-\delta+X(b_t,E_{j^*})\big) + V_{k^*}^{(\sigma,\mu)}[h_t,b_t] \\ &\quad + (c_{k^*}-c_{j^*})A_{k^*}^{(\sigma,\mu)}[h_t,b_t], \end{aligned} \tag{8}$$

where the first line follows from the fact that type c_{j^*} is the highest type in the support of the principal's beliefs in period $t+1$, so he receives a payoff of 0 from $t+1$ onward; the second line follows from part (i) and Lemma A.1, which imply that type c_{j^*} takes the action in periods $t' \geq t+1$ if and only if $b_{t'} \in E_{j^*}$ (note that type c_{j^*} also takes the action at time t); and the third inequality follows from (6).

On the other hand, by Lemma 0(ii), the payoff that type c_{i^*} gets by rejecting the offer at time t is equal to the payoff she would get by mimicking type c_{k^*}, since the principal will believe for sure that the agent's type is not in $\{c_{i^*+1},\dots,c_{j^*}\} \subseteq C[h_t]$ after observing a rejection. That is, type c_{i^*}'s payoff is

$$V_{i^*}^{(\sigma,\mu)}[h_t,b_t] = V_{i^*\to k^*}^{(\sigma,\mu)}[h_t,b_t] = V_{k^*}^{(\sigma,\mu)}[h_t,b_t] + (c_{k^*}-c_{i^*})A_{k^*}^{(\sigma,\mu)}[h_t,b_t]. \tag{9}$$

From equations (8) and (9), it follows that

$$V_{i^*}^{(\sigma,\mu)}[h_t,b_t] - V_{i^*\to j^*}^{(\sigma,\mu)}[h_t,b_t] \leq (c_{j^*}-c_{i^*})\big[A_{k^*}^{(\sigma,\mu)}[h_t,b_t] - \big[1-\delta+X(b_t,E_{j^*})\big]\big] < 0,$$

where the strict inequality follows after using (7). Hence, type c_{i^*} strictly prefers to mimic type c_{j^*} and take the action at time t than to not take it, a contradiction. Hence, all types $c_i \in C[h_t]$ with $c_i \leq c_{j^*}$ take the action at t, and so the set of types taking the action takes the form $C^- = \{c_j \in C[h_t] : c_j < c_{k^*}\}$.

Finally, it is clear that in equilibrium, the transfer that the principal will pay at time t if all agents with type $c_i \in C^-$ take the action is given by ($*$) in Theorem 1. The payoff that an agent with type $c_{j^*} = \max C^-$ gets by accepting the offer is $(1-\delta)(T_t - c_{j^*})$, while her payoff from rejecting the offer and mimicking type $c_{k^*} = \min C[h_t] \setminus C^-$ is $V_{k^*}^{(\sigma,\mu)}[h_t,b_t] + (c_{k^*}-c_{j^*})A_{k^*}^{(\sigma,\mu)}[h_t,b_t]$. Hence, the lowest offer that a c_{j^*} agent accepts is $(1-\delta)T_t = (1-\delta)c_{j^*} + V_{k^*}^{(\sigma,\mu)}[h_t,b_t] + (c_{k^*}-c_{j^*})A_{k^*}^{(\sigma,\mu)}[h_t,b_t]$. *Q.E.D.*

A.1.3. *Proof of Existence and Uniqueness*

For each history h_t and each $c_j \in C[h_t]$, let $C_{j+}[h_t] = \{c_i \in C[h_t] : c_i \geq c_j\}$. For each history h_t and each state realization $b_t \in \mathcal{B}$, let

$$A_{j+}^{(\sigma,\mu)}[h_t,b_t] := (1-\delta)\mathbb{E}_{C_{j+}[h_t]}^{(\sigma,\mu)}\left[\sum_{t'=t+1}^{\infty} \delta^{t'-t} a_{t',j} \,\middle|\, (h_t,b_t)\right],$$

where $\mathbb{E}_{C_{j+}[h_t]}^{(\sigma,\mu)}[\cdot]$ is the expectation operator over the equilibrium strategies when the beliefs of the principal at time $t+1$ have support $C_{j+}[h_t]$. That is, $A_{j+}^{(\sigma,\mu)}[h_t,b_t]$ is the expected discounted fraction of time that an agent with type c_j takes the action after history (h_t,b_t) if the beliefs of the principal at time $t+1$ have support $C_{j+}[h_t]$.

LEMMA A.2: *Fix any equilibrium* (σ, μ) *and history* (h_t, b_t). *The following statements are equivalent*:

(i) *There exists an offer* $T \geq 0$ *such that types* $c_i \in C[h_t]$, $c_i < c_j$, *accept at time* t *and types* $c_i \in C[h_t]$, $c_i \geq c_j$, *reject*.

(ii) *We have* $A_{j+}^{(\sigma,\mu)}[h_t, b_t] \leq 1 - \delta$.

PROOF: First, suppose such an offer T exists and let c_k be the highest type in $C[h_t]$ that accepts T. Let c_j be the lowest type in $C[h_t]$ that rejects the offer and note that $c_k < c_j$. By Lemma 0, the expected discounted payoff that an agent with type c_k gets from accepting the offer is $(1-\delta)(T-c_k) + \delta 0$. The payoff that type c_k obtains by rejecting the offer and mimicking type c_j from time $t+1$ onward is $V_j^{(\sigma,\mu)}[h_t, b_t] + (c_j - c_k)A_{j+}^{(\sigma,\mu)}[h_t, b_t]$. Therefore, the offer T that the principal makes must satisfy

$$(1-\delta)(T-c_k) \geq V_j^{(\sigma,\mu)}[h_t, b_t] + (c_j - c_k)A_{j+}^{(\sigma,\mu)}[h_t, b_t]. \tag{10}$$

Note that an agent with type c_j can guarantee himself a payoff of $(1-\delta)(T-c_j)$ by taking the action in period t and then never taking it again; therefore, incentive compatibility implies

$$\begin{aligned} &V_j^{(\sigma,\mu)}[h_t, b_t] \geq (1-\delta)(T-c_j) \geq V_j^{(\sigma,\mu)}[h_t, b_t] + (c_j - c_k)\big[A_{j+}^{(\sigma,\mu)}[h_t, b_t] - (1-\delta)\big] \\ &\quad \Longrightarrow \quad 1-\delta \geq A_{j+}^{(\sigma,\mu)}[h_t, b_t], \end{aligned}$$

where the second inequality in the first line follows after substituting T from (10).

Suppose next that $A_{j+}^{(\sigma,\mu)}[h_t, b_t] \leq 1-\delta$ and suppose the principal makes offer T such that $(1-\delta)(T-c_k) = V_j^{(\sigma,\mu)}[h_t, b_t] + (c_j - c_k)A_{j+}^{(\sigma,\mu)}[h_t, b_t]$, which only agents with type $c_\ell \in C[h_t]$, $c_\ell \leq c_k$ are supposed to accept. The payoff that an agent with cost c_k obtains by accepting the offer is $(1-\delta)(T-c_k)$, which is exactly what he would obtain by rejecting the offer and mimicking type c_j. Hence, type c_k has an incentive to accept such an offer. Similarly, one can check that all types $c_\ell \in C[h_t]$, $c_\ell < c_k$, also have an incentive to accept the offer. If the agent accepts such an offer and takes the action in period t, the principal will believe that the agent's type lies in $\{c_\ell \in C[h_t] : c_\ell \leq c_i\}$. Note that in all periods $t' > t$, the principal will never offer $T_{t'} > c_k$.

Consider the incentives of an agent with type $c_i \geq c_j > c_k$ at time t. The payoff that this agent gets from accepting the offer is $(1-\delta)(T-c_i)$, since from $t+1$ onward the agent will never accept any equilibrium offer. This is because all subsequent offers will be lower than $c_k < c_j \leq c_i$. On the other hand, the agent's payoff from rejecting the offer is

$$\begin{aligned} V_i^{(\sigma,\mu)}[h_t, b_t] &\geq V_{i\to j}^{(\sigma,\mu)}[h_t, b_t] = V_j^{(\sigma,\mu)}[h_t, b_t] + (c_j - c_i)A_{j+}^{(\sigma,\mu)}[h_t, b_t] \\ &\geq (1-\delta)(T-c_i) = (1-\delta)(c_k - c_i) + V_j^{(\sigma,\mu)}[h_t, b_t] + (c_j - c_k)A_{j+}^{(\sigma,\mu)}[h_t, b_t], \end{aligned}$$

where the second inequality follows since $A_{j+}^{(\sigma,\mu)}[h_t, b_t] \leq 1-\delta$. *Q.E.D.*

The proof of existence and uniqueness relies on Lemma A.2 and uses strong induction on the cardinality of $C[h_t]$. Clearly, Σ_1 is nonempty, and all PBE in Σ_1 give the same payoff to the principal at histories (h_t, b_t) such that $C[h_t] = \{c_k\}$: in this case, the principal offers the agent a transfer $T_{t'} = c_k$ (which the agent accepts) at times $t' \geq t$ such that $b_{t'} \in E_k$ and offers some transfer $T_{t'} < c_k$ (which the agent rejects) at times $t' \geq t$ such that $b_{t'} \notin E_k$.

Suppose next that Σ_{k-1} is nonempty for all $k \leq n-1$ and that for all $k \leq n-1$, all PBE in Σ_k give the principal the same payoff at histories (h_t, b_t) with $|C[h_t]| = k$. We now show that Σ_n is nonempty and that all PBEs in Σ_n give the principal the same payoff at histories (h_t, b_t) with $|C[h_t]| = n$.

Consider a history (h_t, b_t) with $|C[h_t]| = n$. If $b_t \in E_{\overline{k}[h_t]}$, then by part (i), it must be that all agent types in $C[h_t]$ take the action in period t and $T_t = c_{\overline{k}[h_t]}$; hence, at such histories

$$U^{(\sigma,\mu)}[h_t, b_t] = (1-\delta)(b_t - c_{\overline{k}[h_t]}) + \delta \mathbb{E}\big[U^{(\sigma,\mu)}[h_{t+1}, b_{t+1}]|b_t\big].$$

If $b_t \notin E_{\overline{k}[h_t]}$ and $X(b_t, E_{\overline{k}[h_t]}) > 1-\delta$, then by part (ii), all agent types in $C[h_t]$ do not take the action (in this case, the principal makes an offer T small enough that all agents reject); hence, at such states

$$U^{(\sigma,\mu)}[h_t, b_t] = \delta \mathbb{E}\big[U^{(\sigma,\mu)}[h_{t+1}, b_{t+1}]|b_t\big].$$

In either case, the principal does not learn anything about the agent's type, since all types of agents in $C[h_t]$ take the same action, so her beliefs do not change.

Finally, consider states $b_t \notin E_{\overline{k}[h_t]}$ with $X(b_t, E_{\overline{k}[h_t]}) \leq 1-\delta$. Two things can happen at such a state: (a) all types of agents in $C[h_t]$ do not take the action or (b) a strict subset of types in $C[h_t]$ do not take the action and the rest do.[17] In case (a), the principal's beliefs at time $t+1$ would be the same as her beliefs at time t, and her payoffs are

$$U^{(\sigma,\mu)}[h_t, b_t] = \delta \mathbb{E}\big[U^{(\sigma,\mu)}[h_{t+1}, b_{t+1}]|b_t\big].$$

In case (b), the set of types of the agent not taking the action has the form $C_{j+}[h_t] = \{c_i \in C[h_t] : c_i \geq c_j\}$ for some $c_j \in C[h_t]$. So in case (b), the support of the beliefs of the principal at time $t+1$ would be $C_{j+}[h_t]$ if the agent does not take the action and $C[h_t] \setminus C_{j+}[h_t]$ if he does.

By Lemma A.2, there exists an offer that types $C_{j+}[h_t]$ reject and types $C[h_t] \setminus C_{j+}[h_t]$ accept if and only if $A^{(\sigma,\mu)}_{j+}[h_t, b_t] \leq 1-\delta$. Note that, by the induction hypothesis, $A^{(\sigma,\mu)}_{j+}[h_t, b_t]$ is uniquely determined.[18] Let $C^*[h_t, b_t] = \{c_i \in C[h_t] : A^{(\sigma,\mu)}_{i+}[h_t, b_t] \leq 1-\delta\}$. Without loss of generality, renumber the types in $C[h_t]$ so that $C[h_t] = \{c_1, \dots, c_{\overline{k}[h_t]}\}$, with $c_1 < \cdots < c_{\overline{k}[h_t]}$. For each $c_i \in C^*[h_t, b_t]$, let

$$T^*_{t,i-1} = c_{i-1} + \frac{1}{1-\delta}\big(V^{(\sigma,\mu)}_i[h_t, b_t] + A^{(\sigma,\mu)}_{i+}[h_t, b_t](c_i - c_{i-1})\big)$$

be the offer that leaves an agent with type c_{i-1} indifferent between accepting and rejecting when all types in $C_{i+}[h_t]$ reject the offer and all types in $C[h_t] \setminus C_{i+}[h_t]$ accept. Note that $T^*_{t,i-1}$ is the best offer for a principal who wants to get all agents with types in $C[h_t] \setminus C_{i+}[h_t]$ to take the action and all agents with types in types in $C_{i+}[h_t]$ to not take the action.

Let $\mathcal{T} = \{T^*_{t,i-1} : c_i \in C^*[h_t, b_t]\}$. At histories (h_t, b_t) with $b_t \notin E_{\overline{k}[h_t]}$ and with $X(b_t, E_{\overline{k}[h_t]}) \leq 1-\delta$, the principal must choose optimally whether to make an offer in $\mathcal{T}$ or to make a low offer (for example, $T_t = 0$) that all agents reject: an offer $T_t = T^*_{t,i-1}$ would be accepted by types in $C[h_t] \setminus C_{i+}[h_t]$ and rejected by types in $C_{i+}[h_t]$, while an offer $T_t = 0$

[17]By Lemma A.1, in equilibrium an agent with cost $c_{\overline{k}[h_t]}$ does not take the action.

[18]The term $A^{(\sigma,\mu)}_{j+}[h_t, b_t]$ is determined in equilibrium when the principal has beliefs with support $C_{j+}[h_t]$, and the induction hypothesis states that the continuation equilibrium is unique when the cardinality of the support of the principal's beliefs is less than n.

will be rejected by all types. For each offer $T^*_{t,i-1} \in \mathcal{T}$, let $p(T^*_{t,i-1})$ be the probability that offer $T^*_{t,i-1}$ is accepted; that is, the probability that the agent has cost weakly smaller than c_{i-1}. Let $U^{(\sigma,\mu)}[h_t, b_t, T^*_{t,i-1}, a_t = 1]$ and $U^{(\sigma,\mu)}[h_t, b_t, T^*_{t,i-1}, a_t = 0]$ denote the principal's expected continuation payoffs if the offer $T^*_{t,i-1} \in \mathcal{T}$ is accepted and rejected, respectively, at history (h_t, b_t). Note that these payoffs are uniquely pinned down by the induction hypothesis: after observing whether the agent accepted or rejected the offer, the cardinality of the support of the principal's beliefs will be weakly lower than $n-1$. For all (h_t, b_t) with $X(b_t, E_{\overline{k}[h_t]}) \leq 1-\delta$, let

$$U^*(h_t, b_t) = \max_{T \in \mathcal{T}} \bigl\{ p(T)\bigl((1-\delta)(b-T) + U^{(\sigma,\mu)}[h_t, b_t, T, 1]\bigr) \\ + \bigl(1-p(T)\bigr)U^{(\sigma,\mu)}[h_t, b_t, T, 0]\bigr\}$$

and let $T(b_t)$ be a maximizer of this expression.

Partition the states $\mathcal{B}$ as

$$B_1 = E_{\overline{k}[h_t]},$$
$$B_2 = \bigl\{b \in \mathcal{B} \setminus B_1 : X(b, E_{\overline{k}[h_t]}) > 1-\delta\bigr\},$$
$$B_3 = \bigl\{b \in \mathcal{B} \setminus B_1 : X(b, E_{\overline{k}[h_t]}) \leq 1-\delta\bigr\}.$$

By our arguments above, the principal's payoff $U^{(\sigma,\mu)}[h_t, b_t]$ satisfies

$$U^{(\sigma,\mu)}[h_t, b_t] = \begin{cases} (1-\delta)(b_t - c_{\overline{k}[h_t]}) + \delta\mathbb{E}\bigl[U^{(\sigma,\mu)}[h_{t+1}, b_{t+1}]|b_t\bigr], & \text{if } b_t \in B_1, \\ \delta\mathbb{E}\bigl[U^{(\sigma,\mu)}[h_{t+1}, b_{t+1}]|b_t\bigr], & \text{if } b_t \in B_2, \\ \max\bigl\{U^*(h_t, b_t), \delta\mathbb{E}\bigl[U^{(\sigma,\mu)}[h_{t+1}, b_{t+1}]|b_t\bigr]\bigr\}, & \text{if } b_t \in B_3. \end{cases} \tag{11}$$

Let $\mathcal{F}$ be the set of functions from $\mathcal{B}$ to $\mathbb{R}$ and let $\Phi : \mathcal{F} \to \mathcal{F}$ be the operator such that, for every $f \in \mathcal{F}$,

$$\Phi(f)(b) = \begin{cases} (1-\delta)(b - c_{\overline{k}[h_t]}) + \delta\mathbb{E}\bigl[f[b_{t+1}]|b_t = b\bigr], & \text{if } b \in B_1, \\ \delta\mathbb{E}\bigl[f[b_{t+1}]|b_t = b\bigr], & \text{if } b \in B_2, \\ \max\bigl\{U^*(h_t, b), \delta\mathbb{E}\bigl[f[b_{t+1}]|b_t = b\bigr]\bigr\}, & \text{if } b \in B_3. \end{cases}$$

One can check that Φ is a contraction of modulus $\delta < 1$ and, therefore, has a unique fixed point. Moreover, by (11), the principal's equilibrium payoffs $U^{(\sigma,\mu)}[h_t, b_t]$ are a fixed point of Φ. These two observations together imply that the principal's equilibrium payoffs $U^{(\sigma,\mu)}[h_t, b_t]$ are unique. The equilibrium strategies at (h_t, b_t) can be immediately derived from (11). Finally, it can be readily seen that these equilibrium strategies can be taken to be Markovian with respect to the principal's beliefs $\mu[h_t]$ and the shock b_t. *Q.E.D.*

A.2. *Proof of Proposition 3*

Fix a history h_t such that $|C[h_t]| \geq 2$ and without loss of generality renumber the types so that $C[h_t] = \{c_1, \dots, c_{\overline{k}[h_t]}\}$ with $c_1 < \cdots < c_{\overline{k}[h_t]}$. We start by showing that for every such history, there exists a shock realization $b \in \mathcal{B}$ with the property that if $b_s = b$ at time $s \geq t$, then the principal makes an offer that a strict subset of the types in $C[h_t]$ accept.

Suppose for the sake of contradiction that this is not true. Note that this implies that $\mu[h_{t'}] = \mu[h_t]$ for every $h_{t'} \succeq h_t$. By Theorem 1, this further implies that after history h_t, the agent only takes the action when the shock is in $E_{\overline{k}[h_t]}$ and receives a transfer equal to $c_{\overline{k}[h_t]}$. Therefore, the principal's payoff after history (h_t, b_t) is

$$U^{(\sigma,\mu)}[h_t, b_t] = (1-\delta)\mathbb{E}\left[\sum_{t'=t}^{\infty} \delta^{t'-t}(b_{t'} - c_{\overline{k}[h_t]})\mathbf{1}_{\{b_{t'} \in E_{\overline{k}[h_t]}\}} \middle| b_t = b\right].$$

Let $b \in E_{\overline{k}[h_t]-1}$ be such that $X(b, E_{\overline{k}[h_t]}) < 1-\delta$. The conditions in the statement of Proposition 3 guarantee that such a shock b exists. Suppose that the shock at time $s \geq t$ is $b_s = b$ and let $\varepsilon > 0$ be small enough such that

$$T = c_{\overline{k}[h_t]-1} + \frac{1}{1-\delta} X(b, E_{\overline{k}[h_t]})(c_{\overline{k}[h_t]} - c_{\overline{k}[h_t]-1}) + \varepsilon < c_{\overline{k}[h_t]}. \tag{12}$$

Note that at history (h_s, b_s), an offer equal to T is accepted by all types with cost strictly lower than $c_{\overline{k}[h_t]}$ and is rejected by type $c_{\overline{k}[h_t]}$.[19] The principal's payoff from making an offer T conditional on the agent's type being $c_{\overline{k}[h_t]}$ is $U^{(\sigma,\mu)}[h_t, b_t]$. On the other hand, when the agent's type is lower than $c_{\overline{k}[h_t]}$, the principal obtains $(1-\delta)(b-T)$ at period t if she offers transfer T, and learns that the agent's type is not $c_{\overline{k}[h_t]}$. From period $t+1$ onward, the principal's payoff is bounded below by what she could obtain if at all periods $t' > t$ she offers $T_{t'} = c_{\overline{k}[h_t]-1}$ whenever $b_{t'} \in E_{\overline{k}[h_t]-1}$ (an offer that is accepted by all types) and offers $T_{t'} = 0$ otherwise (that is rejected by all types). The payoff that the principal obtains from following this strategy when the agent's cost is lower than $c_{\overline{k}[h_t]}$ is

$$\begin{aligned}
\underline{U} &= (1-\delta)(b-T) + (1-\delta)\mathbb{E}\left[\sum_{t'=s+1}^{\infty} \delta^{t'-s}(b_{t'} - c_{\overline{k}[h_t]-1})\mathbf{1}_{\{b_{t'} \in E_{\overline{k}[h_t]-1}\}} \middle| b_s = b\right] \\
&= (1-\delta)(b - c_{\overline{k}[h_t]-1} - \varepsilon) + (1-\delta)\mathbb{E}\left[\sum_{t'=s+1}^{\infty} \delta^{t'-s}(b_{t'} - c_{\overline{k}[h_t]})\mathbf{1}_{\{b_{t'} \in E_{\overline{k}[h_t]}\}} \middle| b_s = b\right] \\
&\quad + (1-\delta)\mathbb{E}\left[\sum_{t'=s+1}^{\infty} \delta^{t'-s}(b_{t'} - c_{\overline{k}[h_t]-1})\mathbf{1}_{\{b_{t'} \in E_{\overline{k}[h_t]-1} \setminus E_{\overline{k}[h_t]}\}} \middle| b_s = b\right] \\
&= U^{(\sigma,\mu)}[h_t, b] + (1-\delta)(b - c_{\overline{k}[h_t]-1} - \varepsilon) \\
&\quad + (1-\delta)\mathbb{E}\left[\sum_{t'=s+1}^{\infty} \delta^{t'-s}(b_{t'} - c_{\overline{k}[h_t]-1})\mathbf{1}_{\{b_{t'} \in E_{\overline{k}[h_t]-1} \setminus E_{\overline{k}[h_t]}\}} \middle| b_s = b\right],
\end{aligned}$$

where the second line follows from substituting (12). Since $b \in E_{\overline{k}[h_t]-1}$, from the third line it follows that if $\varepsilon > 0$ is small enough, then $\underline{U}$ is strictly larger than $U^{(\sigma,\mu)}[h_t, b]$. But this cannot be, since the proposed strategy profile was an equilibrium. Therefore, for all histories h_t such that $|C[h_t]| \geq 2$, there exists $b \in \mathcal{B}$ with the property that at history (h_s, b_s) with $h_s \succeq h_t$ and $b_s = b$, the principal makes an offer that a strict subset of the types in $C[h_t]$ accept.

[19] By accepting offer T, an agent with cost $c_i < c_{\overline{k}[h_t]}$ obtains a payoff of at least $(1-\delta)(T-c_i) + \delta 0$. This payoff is strictly larger than the payoff of $X(b, E_{\overline{k}[h_t]})(c_{\overline{k}[h_t]} - c_i)$ he obtains by rejecting and continuing to play the equilibrium.

We now use this result to establish the proposition. Note first that this result, together with the assumption that process $\{b_t\}$ is ergodic, implies that there is *long-run learning* in equilibrium. This is because as long as $C[h_t]$ has two or more elements, there will be some shock realization at which the principal makes an offer that only a strict subset of types in $C[h_t]$ accept. And since there are finitely many types and $\{b_t\}$ is ergodic, it is certain that the principal will end up learning the agent's type.

Finally, fix a history h_t such that $C[h_t] = \{c_i\}$. Then, from time t onward the principal's payoff is $U^{(\sigma,\mu)}[h_t, b] = (1-\delta)\mathbb{E}[\sum_{t'=t}^{\infty} \delta^{t'-t}(b_{t'} - c_i)\mathbf{1}_{\{b_{t'} \in E_i\}} | b_t = b] = U_i^*(b|c = c_i)$, which is the first-best payoff. This and the previous arguments imply that the equilibrium is long-run first-best. *Q.E.D.*

REFERENCES

ACHARYA, A., AND J. ORTNER (2017): "Supplement to 'Progressive Learning,'" *Econometrica Supplemental Material*, 85, http://dx.doi.org/10.3982/ECTA14718. [1970]

BLUME, A. (1998): "Contract Renegotiation With Time-Varying Valuations," *Journal of Economics & Management Strategy*, 7, 397–433. [1967]

BOARD, S., AND M. PYCIA (2014): "Outside Options and the Failure of the Coase Conjecture," *American Economic Review*, 104, 656–671. [1967]

CALLANDER, S., AND N. MATOUSCHEK (2014): "Managing on Rugged Landscapes," Technical Report, Northwestern University. [1967]

CARMICHAEL, H. L., AND W. B. MACLEOD (2000): "Worker Cooperation and the Ratchet Effect," *Journal of Labor Economics*, 18, 1–19. [1966]

CHASSANG, S. (2010): "Building Routines: Learning, Cooperation, and the Dynamics of Incomplete Relational Contracts," *American Economic Review*, 100, 448–465. [1967]

COMPTE, O., AND P. JEHIEL (2002): "On the Role of Outside Options in Bargaining With Obstinate Parties," *Econometrica*, 70, 1477–1517. [1967]

DEWATRIPONT, M. (1989): "Renegotiation and Information Revelation Over Time: The Case of Optimal Labor Contracts," *Quarterly Journal of Economics*, 104, 589–619. [1965]

DILLEN, M., AND M. LUNDHOLM (1996): "Dynamic Income Taxation, Redistribution, and the Ratchet Effect," *Journal of Public Economics*, 59, 69–93. [1965]

FIOCCO, R., AND R. STRAUSZ (2015): "Consumer Standards as a Strategic Device to Mitigate Ratchet Effects in Dynamic Regulation," *Journal of Economics & Management Strategy*, 24, 550–569. [1966]

FREIXAS, X., R. GUESNERIE, AND J. TIROLE (1985): "Planning Under Incomplete Information and the Ratchet Effect," *Review of Economic Studies*, 52, 173–191. [1965]

FUCHS, W., AND A. SKRZYPACZ (2010): "Bargaining With Arrival of New Traders," *American Economic Review*, 100, 802–836. [1967]

FUDENBERG, D., D. K. LEVINE, AND J. TIROLE (1985): "Infinite-Horizon Models of Bargaining With One-Sided Incomplete Information," in *Bargaining With Incomplete Information*, ed. by A. Roth. Cambridge: Cambridge University Press, 73–98. [1967]

GERARDI, D., AND L. MAESTRI (2015): "Dynamic Contracting With Limited Commitment and the Ratchet Effect," Technical Report, Collegio Carlo Alberto. [1967,1971]

GIBBONS, R. (1987): "Piece-Rate Incentive Schemes," *Journal of Labor Economics*, 5, 413–429. [1965]

GIBBONS, R., AND R. HENDERSON (2012): "What Do Managers Do?: Exploring Persistent Performance Differences Among Seemingly Similar Enterprises," Working Paper, Harvard Business School. [1967]

GUL, F., H. SONNENSCHEIN, AND R. WILSON (1986): "Foundations of Dynamic Monopoly and the Coase Conjecture," *Journal of Economic Theory*, 39, 155–190. [1967]

HALAC, M. (2012): "Relational Contracts and the Value of Relationships," *American Economic Review*, 102, 750–779. [1965]

HALAC, M., AND A. PRAT (2016): "Managerial Attention and Worker Performance," *American Economic Review*, 106, 3104–3132. [1967]

HART, O. D., AND J. TIROLE (1988): "Contract Renegotiation and Coasian Dynamics," *Review of Economic Studies*, 55, 509–540. [1965-1967,1969-1971,1973]

HART, S. (1985): "Nonzero-sum Two-Person Repeated Games With Incomplete Information," *Mathematics of Operations Research*, 10, 117–153. [1968]

HWANG, I., AND F. LI (2017): "Transparency of Outside Options in Bargaining," *Journal of Economic Theory*, 167, 116–147. [1967]

KAHN, L. B. (2010): "The Long-Term Labor Market Consequences of Graduating From College in a Bad Economy," *Labour Economics*, 17, 303–316. [1966,1977]

KANEMOTO, Y., AND W. B. MACLEOD (1992): "The Ratchet Effect and the Market for Secondhand Workers," *Journal of Labor Economics*, 10, 85–98. [1966]

KENNAN, J. (2001): "Repeated Bargaining With Persistent Private Information," *Review of Economic Studies*, 68, 719–755. [1967]

LAFFONT, J.-J., AND J. TIROLE (1988): "The Dynamics of Incentive Contracts," *Econometrica*, 56, 1153–1175. [1965]

LAZEAR, E. P., K. L. SHAW, AND C. STANTON (2016): "Making Do With Less: Working Harder During Recessions," *Journal of Labor Economics*, 34, S333–S360. [1966,1976]

LI, J., AND N. MATOUSCHEK (2013): "Managing Conflicts in Relational Contracts," *American Economic Review*, 103, 2328–2351. [1967]

MALCOMSON, J. M. (2016): "Relational Incentive Contracts With Persistent Private Information," *Econometrica*, 84, 317–346. [1965]

OREOPOULOS, P., T. VON WACHTER, AND A. HEISZ (2012): "The Short- and Long-Term Career Effects of Graduating in a Recession," *American Economic Journal: Applied Economics*, 4, 1–29. [1966,1977]

ORTNER, J. (2017): "Durable Goods Monopoly With Stochastic Costs," *Theoretical Economics*, 12, 817–861. [1967]

PESKI, M. (2008): "Repeated Games With Incomplete Information on One Side," *Theoretical Economics*, 3, 29–84. [1968,1974]

——— (2014): "Repeated Games With Incomplete Information and Discounting," *Theoretical Economics*, 9, 651–694. [1968]

SCHMIDT, K. M. (1993): "Commitment Through Incomplete Information in a Simple Repeated Bargaining Game," *Journal of Economic Theory*, 60, 114–139. [1965,1966,1969-1971,1973]

SORIN, S. (1999): "Merging, Reputation, and Repeated Games With Incomplete Information," *Games and Economic Behavior*, 29, 274–308. [1968]

WATSON, J. (1999): "Starting Small and Renegotiation," *Journal of Economic Theory*, 85, 52–90. [1968]

——— (2002): "Starting Small and Commitment" *Games and Economic Behavior* 38, 176–199. [1968]

WISEMAN, T. (2005): "A Partial Folk Theorem for Games With Unknown Payoff Distributions," *Econometrica*, 73, 629–645. [1968]

Co-editor Dirk Bergemann handled this manuscript.

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