Reputation and Earnings Dynamics^{*}

Boyan Jovanovic[†] and Julien Prat[‡]

October 2020

(First version: September 2016)

Abstract

Cyclical patterns in earnings can arise when contracts between firms and their workers are incomplete, and when workers cannot borrow or lend so as to smooth their consumption. Effort cycles generate occasional large changes in earnings. These large changes are transitory, consistent with recent empirical findings.

Keywords: Reputation, Career Concerns, Earnings Dynamics, Endogenous Fluctuations.

1 Introduction

Recent analyses of large panel data sets reveal that earnings fluctuate in ways that cannot be captured by the linear models commonly used in the literature on earnings dynamics. In particular, earnings data show occasional changes far larger than predicted by a log-normal distribution – see Guvenen, Karahan, Ozkan and Song (2015) and Arellano, Blundell and Bonhomme (2017), henceforth GKOS and ABB respectively.

^{*}Previously circulated under the title "Reputation Cycles".

[†]Economics Department, NYU. bj2@nyu.edu

[‡]CNRS, CREST, Ecole Polytechnique. julien.prat@ensae.fr. We are grateful to Sai Ma and Har-gunghet Singh for excellent research assistance and to Jess Benhabib, Emmanuel Farhi, Jean-Michel Grandmont, Fatih Guvenen, Dimitris Papanikolaou, David Pearce and Gianluca Violante for their insightful suggestions. We thank Fatih Guvenen for providing us with data and participants of the SED 2015, Barcelona GSE summer forum 2015, NASM 2016, and Minnesota Macro 2016 for their comments. Boyan Jovanovic thanks the NSF for its 2015-2018 grant for the project "Reputations in Macroeconomics." Julien Prat acknowledges the support of the Investissements d'Avenir grant (ANR-11-IDEX-0003/Labex Ecodec/ANR-11-LABX-0047).

These unusual events are routinely attributed to external shocks such as layoffs triggering falls off the job ladder. Since large income fluctuations reduce the welfare of risk-averse agents, it seems natural to conjecture that they are not initiated by them. We show that career-concerned agents may in fact prefer to go through periods of low and high effort. We call these oscillations *reputation cycles*.

We change the career concerns model of Holmström (1999) by assuming that the agent is risk-averse and cannot borrow or lend. Since reputational investment pays off only in the future, a worker's incentive to maintain or improve her reputation depends on her discount factor. If she is risk averse, her discount factor depends negatively on her consumption growth. When the worker cannot smooth her consumption by other means,¹ this force can then give rise to cycles in her effort, and thus in her earnings. The mechanism works as follows: When effort is low, the worker's consumption is low relative to her future consumption which means that her discount factor is also low. This reduces her incentive to create a good reputation and low effort is self fulfilling. In other words, the discount factor and reputational concerns are low because consumption growth is high. Conversely, in a period of high effort, current consumption is higher than future consumption, which means that the discount factor is also large and so high effort is self fulfilling too.

The model delivers several insights. First, evidence has shown that tangible investment is inhibited by liquidity constraints – Fazzari, Hubbard and Petersen (1988). We find that the opposite may hold, at least as far as reputational investment – an intangible – is concerned. Under the set of calibrated parameters, cyclical paths may entail higher average reputation capital than does the rest point solution of the model, although cyclical paths exist only if the agent cannot borrow or lend.

Second, in terms of pure theory, we establish the existence of a new class of equilibria. Repeated games are known to have equilibria that shift from period to period with no change in the payoff functions in the stage game – e.g., Green and Porter (1984), Mailath and Samuelson (2016, Sec. 4.3.1). In dynamic games with type uncertainty of which ours is an example, Markovian equilibria can also entail fluctuating effort – e.g., Board and Meyer-ter-Vehn (2013) and Cisternas (2018). With a risk averse agent, however, we show that the Holmström model acquires a cyclical equilibrium.²

¹Kaplan, Weidner and Violante (2014) estimate that one third of U.S. consumers are hand to mouth.

²The career concerns literature focuses on risk-neutral agents. A notable exception is Bouvard and Levy (2019) who consider a problem where the sender's payoff function is concave because she is facing heterogeneous receivers. Bouvard and Levy also find that concave preferences generate multiple equilibria even though they do not analyze self-sustaining cycles but focus instead on the emergence

Moreover, cyclical equilibria can arise even though the agent has a publicly known retirement date.

Third, we explore the model's quantitative properties using micro-data on earnings dynamics to calibrate its parameters. The variance coefficients can be recovered from Lange's (2007) estimates about the speed of employer learning. In the calibration of the model with stochastic cycles, income growth decomposes into small, fully persistent changes and larger, mean-reverting changes (see Fig. 7). These predictions are in line with recent evidence from U.S. panel data on earnings: GKOS and ABB detect significant deviations from lognormality, with a small but noticeable share of individuals experiencing very large changes. They also find that the degree of persistence appears to be non-linear in the size of the shocks as large shocks exhibit much stronger mean reversion than small shocks. Hubmer (2018) shows that a life-cycle version of the standard job ladder model can capture the large negative skewness and high excess kurtosis of earnings data. Our analysis indicates that endogenous fluctuations in effort can also help to explain these large deviations from the log-normal framework, while addressing the systematic correlation between the persistence and the size of income shocks.

Finally, the model can generate a *negative* relation between exogenous income noise and the volatility of earnings: Income can become more volatile when reputational concerns rise in response to a reduction in the noise accompanying the relation between effort and output. Agency theory suggests that a reduction in noise raises the incentives to provide effort, and we show that this can trigger a rise in the volatility of earnings. For cycles in effort to exist, incentives to provide effort cannot be too weak. Noise with which the agent's output is observed deters effort and it also obstructs the formation of cycles. This bears on the question of why income inequality has risen – Song, Price, Guvenen, Bloom, von Wachter (2018), henceforth SPGBW. Paradoxically, we find that an increase in inequality could have in part been caused by a *decline* in exogenous volatility. In particular, a decline in the noisiness of technology will strengthen reputational concerns and may raise equilibrium income volatility. Hence what SPGBW measure as a rise in the volatility of observed income may in fact be the result of a decline in exogenous output noise.

The plan of the paper is as follows. Section 2 lays out the model. Section 3 characterizes the rest point of the model where effort remains constant, while Section 4 ex-

of an inefficient steady-state.

plains how and when cycles may emerge. We parametrize the model in Section 5 and explore its quantitative implications. Then we review major hypotheses for endogenous cycles in Section 6.1 so as to highlight the novelty of reputation cycles. Several extensions are discussed in Section 6.2, while Section 7 concludes. All the proofs are relegated to the Appendix.

2 Model

In Holmström (1999) a risk-neutral agent faces a spot market with risk-neutral buyers. The unique equilibrium then entails a monotonic time path of effort. We will show that if one assumes that the agent is risk averse, multiple equilibria arise and in some of them effort follows a 2-period cycle. We assume the following:

- A1. Non-contractible effort and output;
- A2. Risk averse agent facing short-term risk-neutral buyers;
- A3. Agent cannot borrow or lend.

We maintain these assumptions throughout.

Production.—Each period an agent produces output y_t by exerting hidden effort a_t :

$$y_t = \theta_t + a_t + \varepsilon_t. \quad (1)$$

Here $\varepsilon_t \sim N(0, \sigma_\varepsilon^2)$ is an i.i.d. shock. The variable θ is the agent's efficiency; θ is unknown, even to the agent, and the common prior is $\mathcal{N}(\bar{\theta}, \sigma_\theta^2)$.

Earnings.—Contracts in which the period- t payment is contingent on y_t are not feasible. Instead, buyers pay the agent up front and the payment reflects the market belief at the beginning of the period. Thus the agent exerts effort only so as to raise her future income. Output has a persistent effect on prices because it affects the market's belief about the agent's efficiency which fluctuates over time as follows

$$\theta_t = \theta_{t-1} + \nu_t, \text{ with } \nu \sim \mathcal{N}(0, \sigma_\nu^2). \quad (2)$$

Information.—Each period the agent chooses an a from $\mathcal{A}$ while everyone observes y . The agent can also observe $\theta + \varepsilon$ but not its components. The market's participants observe y only, but since they will deduce a from equilibrium, they too

will be able to infer $\theta + \varepsilon$. However, the market and the agent observe different signals whenever the agent deviates and take a different action than the one anticipated by the market. We account for such divergence by using $E^*[\cdot]$ and $E[\cdot]$ to denote the market's and the agent's expectation, respectively.

Timing.—Within a period, events unfold as follows:

1. There are many identical and risk neutral potential buyers, and they get zero rents. Hence the up-front payment the agent gets is $E^*[y_t | y^t]$, where $y^t \equiv \{y_s\}_{s=0}^{t-1}$ is her public history.
2. The agent chooses a_t , privately.
3. Output y_t is realized, and it becomes part of the publicly observed history y^{t+1} .

Preferences.—The agent is risk averse and infinitely lived. Her preferences as of date 0 read

$$\sum_{t=0}^{\infty} \beta^t \mathbf{U}(c_t, a_t),$$

with $\beta \in (0, 1)$ denoting the discount factor. To neutralize the wealth effect, we will focus on per-period utility functions that exhibit constant absolute risk aversion so that $\mathbf{U}(c, a) = -\exp(-\gamma(c - g(a)))$.

Equilibrium.—We will examine equilibria in pure strategies in which the agent's behavior along the path of play might depend on calendar time only but not on the history of output.³ An equilibrium must satisfy the following three requirements. First, payments are set competitively by the market. Second, the agent chooses her effort so as to maximize expected payoffs. Third, the market conjecture a_t^* is verified.

Definition 1 A *deterministic equilibrium* is a profile $\{a_t, a_t^*, c_t\}_{t=0}^{\infty}$ such that $\{a_t\}_{t=0}^{\infty}$ is a function of time only and:

1. Expected profits are zero in every period as $c_t = E^*[y_t | y^t] = E^*[\theta_t + a_t | y^t]$.
2. Effort a_t is chosen optimally

$$a_t = \arg \max_{\hat{a}_t \in \mathcal{A}} E_t \left[\mathbf{U}(c_t, \hat{a}_t) + \sum_{s=1}^{\infty} \beta^{t+s} \mathbf{U}(c_{t+s}, a_{t+s}) \right], \text{ for all } t = 0, 1, \dots \quad (3)$$

³Deterministic action profiles will be equilibrium strategies under the CARA-GHH preferences that we shall use – see Proposition 2 and the associated discussion. For GHH preferences see Greenwood, Hercowitz and Huffman (1988).

3. The market's conjecture $\{a_t^*\}_{t=0}^\infty$ is correct, so that $a_t^* = a_t$ for all $t = 0, 1, \dots$

Market beliefs.—In order to interpret the output signal, the market's participants need to conjecture the agent's behavior. For a given conjecture $\{a_s^*\}_{s=0}^{t-1}$, a sufficient statistic for the information revealed about θ_t is the sequence $x^t = \{x_s\}_{s=0}^{t-1}$, where

$$x_t \equiv y_t - a_t^* = \theta_t + a_t - a_t^* + \varepsilon_t. \quad (4)$$

The market treats x as the signal. On path, i.e., when $a_t = a_t^*$, $x_t \sim N(\theta_t, \sigma_\varepsilon^2)$. Because ν and ε are normal, the posterior is also normal: $\theta_t \sim \mathcal{N}(m_t, \sigma_{\theta,t}^2)$. The posterior variance evolves deterministically

$$\sigma_{\theta,t}^2 = \frac{1}{\sigma_{\theta,t-1}^{-2} + \sigma_\varepsilon^{-2}} + \sigma_\nu^2, \quad (5)$$

and converges in the long-run to $\bar{\sigma}_\theta^2$, whose value is obtained setting $\sigma_{\theta,t} = \sigma_{\theta,t-1}$ in (5) so that⁴

$$\bar{\sigma}_\theta^{-2} = \frac{\sigma_\varepsilon^{-2}}{2} \left(\sqrt{1 + 4 \frac{\sigma_\varepsilon^2}{\sigma_\nu^2}} - 1 \right). \quad (6)$$

We will focus on stationary learning (Holmström, 1999, Sec. 2.2). Hence we assume that the agent's initial θ is drawn from $\mathcal{N}(\bar{\theta}, \bar{\sigma}_\theta^2)$. Then $\sigma_{\theta,t} = \bar{\sigma}_\theta$ remain constant and time is not a state. Let m_t denote the mean of the market's belief about θ_t , which we will also refer to as the agent's reputation. m obeys the following law of motion

$$m_{t+1} = E^*[\theta_{t+1} \mid m_t, x_t] = \lambda m_t + (1 - \lambda) x_t, \text{ where } \lambda \equiv \frac{\sigma_\varepsilon^2}{\sigma_\varepsilon^2 + \bar{\sigma}_\theta^2}. \quad (7)$$

As (4) shows, by deviating from a_t^* , the agent can raise x_t , manipulate m_{t+1} and drive a wedge between her own belief and that of the market. We now analyze the agent's incentives to do that.

2.1 Investment in reputation

The zero-profit condition for market's participants is satisfied when $c_t = E^*[y_t \mid y^t] = m_t + a_t^*$, with m_t obeying the law of motion (7). We focus on equilibria in which a_t^* is

⁴Observe that $\partial \bar{\sigma}_\theta^{-2} / \partial \sigma_\varepsilon < 0$ and $\partial \bar{\sigma}_\theta^{-2} / \partial \sigma_\nu < 0$ as fluctuations in output and efficiency lower stationary precision.

deterministic. The agent's action a^* must solve:

$$a_t^* = \arg \max_{a_t \in \mathcal{A}} E_t \left[\sum_{s=0}^{\infty} \beta^{t+s} \mathbf{U} (m_{t+s} + a_{t+s}^*, a_{t+s}) \right], \text{ for all } t = 0, 1, \dots \quad (8)$$

This requirement ensures that: (i) given the law of motion of beliefs derived in the previous subsection, equilibrium strategies satisfying (8) maximize the agent's utility; (ii) given the equilibrium actions, beliefs are updated via Bayes rule which is well defined for all a since all output levels occur with positive probability at a^* . Solutions to problem (8) are therefore optimal.

We now establish a sufficient condition for the existence of deterministic equilibria.⁵ Whenever the agent deviates, she drives a persistent wedge between her belief and that of the market. Thus a deviation that is not attractive on the equilibrium path might nonetheless be profitable off path, and so the necessary condition is not generally sufficient. We now show that this concern does not apply when $U(\cdot)$ is in the CARA-GHH class.

Proposition 1 *Assume that the per-period utility function is*

$$\mathbf{U}(c, a) = U(c - g(a)) = -\exp(-\gamma(c - g(a))), \quad (9)$$

with $\gamma > 0$, $g'(\cdot) > 0$ and $g''(\cdot) \geq 0$. A deterministic profile $\{a_t^\}_{t=0}^{\infty}$ is an equilibrium profile if and only if*

$$-\mathbf{U}_2(c_t, a_t^*) = \frac{1 - \lambda}{\lambda} \sum_{s=t+1}^{\infty} (\beta\lambda)^{s-t} E_t [\mathbf{U}_1(c_s, a_s^*)]. \quad (10)$$

The proof is in the Appendix and it proceeds in two steps. First, we establish that optimal strategies are deterministic on and off the equilibrium path whenever the market's conjecture a_t^* is deterministic.⁶ This result holds because shocks to

⁵Note that the dynamic stability we consider here is quite distinct from the notion of strategic stability explored by Kohlberg and Mertens (1986). The latter, defined for finite games, is concerned with robustness to certain kinds of perturbations of the game in question, while we are looking at stable solutions of a dynamic strategic system.

⁶We consider Markov strategies where the agent's action might depend on the history of the game through calendar time t as well as the market's and the agent's beliefs. Allowing for arbitrary deviations, we show in Lemma 1 that, when the utility function is CARA, optimal strategies on and off the equilibrium path do not depend on the agent's reputation. Hence they are independent of output history and are a function of calendar time only.

the agent's reputation follow a martingale process, a requirement that holds in our stationary learning model as well as in non-stationary environments. Then income shocks, which reflect the history y^t , will in turn drop out of the first-order condition because buyers are risk neutral and their beliefs have the martingale property.⁷ Martingale types of shocks matter for the decisions of the agent only if her preferences display a wealth effect, which can be turned off using the CARA-GHH specification. We can therefore focus, without loss of generality, on sequences of deterministic actions. This restriction enables us to prove the concavity of the agent's objective in the space of action sequences, thus directly establishing the sufficiency of the first order condition.⁸

The left-hand side of the necessary condition (10) is the marginal disutility, while the right-hand side is the discounted benefit because a marginal increase in effort raises the posterior mean at date s by $\partial m_s / \partial a_t = (1 - \lambda) \lambda^{s-t-1}$, for all $s > t$.⁹ Since c is increasing in m and m is a martingale, expected consumption depends on the agent's current reputation. Hence the optimality condition (10) varies with m_t and in general so would a_t^* , contrary to our assumption that a_t^* depends on t only. To ensure that $E_t [\mathbf{U}_1(c_s, a_s^*) / \mathbf{U}_2(c_t, a_t^*)]$ is not affected by m_t , we neutralize the wealth effect by assuming that the agent's utility function is CARA. This restriction ensures that the agent's discount factor no longer depends on m_t because it is the forecastable component of m_s for all $s > t$.

Proposition 2 *When the per-period utility function is CARA, a deterministic profile $\{a_t^*\}_{t=0}^\infty$ is an equilibrium profile with interior effort if and only if*

$$g'(a_t^*) = \frac{1 - \lambda}{\lambda} \sum_{s=t+1}^{\infty} (r\lambda)^{s-t} \exp(\gamma [a_t^* - g(a_t^*) - (a_s^* - g(a_s^*))]), \quad (11)$$

where

$$r \equiv \beta \exp\left(\frac{\gamma^2 \sigma_\nu^2}{2}\right). \quad (12)$$

⁷See the proof of Proposition 2 and, in particular, the derivation of (33) which holds because the current level of m does not carry additional information about its evolution in the future. Note that this would not anymore be the case if the agent's type θ_t were to follow a mean-reverting process instead of a random walk, as stipulated in (2).

⁸A recent paper by Cisternas (2018) provides bounds guaranteeing sufficiency in more general settings in which optimal actions may depend on the posterior m .

⁹Eq. (10) is the counterpart of Holmström's (1999), eq. (22), with the essential addition of the terms $\mathbf{U}_2(c_t, a_t^*)$ and $E_t^* [\mathbf{U}_1(c_s, a_s^*)]$.

The effective discount factor r is higher than β because consumption is stochastic and the agent is prudent ($U'''(\cdot) > 0$). A mean-preserving spread in consumption raises expected marginal utility in future periods, and so makes the agent more patient. This precautionary motive is stronger when the coefficient of risk-aversion γ is large because γ governs the curvature of the marginal utility function. The discount factor r is also increasing in σ_ν^2 because beliefs, and thus consumption, are more volatile when types are less persistent. By contrast, r does not depend on the output variance σ_ε^2 . This surprising result arises because σ_ε^2 has two countervailing effects. To describe them, it is useful to recall that $m_{t+1} = \lambda m_t + (1 - \lambda) x_t$, and that consumption is linear in m as $c_t = m_t + a_t^*$. Hence σ_ε^2 has a positive effect on consumption volatility since it raises the variance of the market signal x . However, σ_ε^2 also raises the gain parameter λ as the market puts less weight on recent observations when production is more noisy. This additional inertia in the updating rule exactly compensates the direct impact on the variance of the signal, thus leaving consumption volatility unchanged.¹⁰ In the remainder of the paper, we guarantee that expected returns are bounded by focusing on cases where $r < 1$.¹¹

3 Rest Point

Before analyzing cycles, we first show that the model has a rest point where effort remains constant over time. Characterizing the stationary equilibrium has two main benefits. First, it will enable us to compare its welfare properties with those of cycles. Second, it will allow us to characterize the region where cycles occur by relating it to the local stability of the rest point. To keep the analysis as general as possible, we distinguish two main configurations: those where the rest point is at a corner of the feasible set because costs are linear in a , and those where the rest point lies in the interior of the feasible set because costs are strictly convex in a .

¹⁰Note that this finding holds because prior precision is set equal to its stationary level $\bar{\sigma}_\theta^{-2}$. For any other value, the conditional variance of m would not anymore be independent of σ_ε^2 .

¹¹Another reason to rule out configurations where $r > 1$ is that they may exhibit excessive levels of effort at the rest point instead of insufficient ones.

3.1 Rest point with linear costs

Suppose that $g(a) = \kappa a$ with $\kappa \in (0, 1)$ to ensure that effort raises output net of disutility. We rule out infinite effort by imposing an upper-bound $\bar{a}$ on the feasible set $\mathcal{A} = [0, \bar{a}]$. With linear costs, optimality can only be restored through changes in the discount factor because both marginal costs and marginal returns are constant. To see why, assume that both a_t^* and a_{t+1}^* belong to the interior of $\mathcal{A}$ and substitute $g'(a_{t+1}^*)$ on the right-hand side of (11) to obtain the recursive condition

$$g'(a_t^*) = \beta E_t^* \left[\frac{U'(c_{t+1} - g(a_{t+1}^*))}{U'(c_t - g(a_t^*))} \right] (1 - \lambda + \lambda g'(a_{t+1}^*)). \quad (13)$$

The recursive incentive constraint highlights that increasing a_t has two benefits:

- (i) It raises $t + 1$ earnings by $(1 - \lambda)$, and
- (ii) it enables the agent to reduce her effort in period $t + 1$ by λ and undo the deviation, thereby restoring the reputation associated to the equilibrium path in $t + 2$ and beyond.

Both benefits are converted into today's utils through multiplication by the stochastic discount factor $\beta E_t [U'(c_{t+1} - g(a_{t+1}^*)) / U'(c_t - g(a_t^*))]$. Since the utility function is CARA, (13) reads

$$g'(a_t^*) = r \exp(\gamma [a_t^* - g(a_t^*) - (a_{t+1}^* - g(a_{t+1}^*))]) (1 - \lambda + \lambda g'(a_{t+1}^*)). \quad (14)$$

Reinserting the linear cost function $g(a) = \kappa a$ into (14), we finally find that

$$\frac{U'(a_{t+1}^*(1 - \kappa))}{U'(a_t^*(1 - \kappa))} = \frac{\kappa}{r(1 + \lambda(\kappa - 1))}. \quad (15)$$

Taking expectations about a_{t+1}^* as given, a_t^* adjusts until the equality is satisfied. The equilibrium path is fully determined by the ratio of marginal costs to discounted marginal returns on the right-hand side of (15). Since this ratio is constant, all paths converge to $\bar{a}$ when it is lower than one or, conversely, converge to zero when it is higher than one. The dynamics become explicit if one uses (9) to rewrite (15) as

$$a_{t+1}^* = a_t^* - \frac{1}{\gamma(1 - \kappa)} \log \left(\frac{\kappa}{r(1 + \lambda(\kappa - 1))} \right). \quad (16)$$

The law of motion is obtained projecting a_t^* on a line parallel to the 45-degree line

with an intercept equal to the constant on the right-hand side of (16). As shown in Appendix A.2, Fig. 10, in the knife-edge case where the intercept is zero, i.e., when $\kappa = (1 - \lambda)/(r^{-1} - \lambda)$, the dynamic map and 45-degree line coincide so that any action in $\mathcal{A}$ is a potential rest point. Besides this particular case, the action path converges to $\bar{a}$ when the intercept is positive or to zero when the intercept is negative. Hence a_t^* can be constant solely at one of the boundary of the action set.

If marginal costs are low, so that $\kappa < (1 - \lambda)/(r^{-1} - \lambda)$, agents are tempted to raise their effort above that of the rest point $\hat{a}$. Hence $\hat{a}$ is sustainable only if such deviations are not feasible, that is if $\hat{a} = \bar{a}$.¹² This solution is also efficient because the cost parameter $\kappa < 1$ while marginal productivity is one. Conversely, when marginal costs are high, so that $\kappa > (1 - \lambda)/(r^{-1} - \lambda)$, reputational concerns are too weak and $\hat{a} = 0$ is the only rest point.

3.2 Rest point with convex costs

We now generalize our setup by assuming that disutility is *strictly convex* in effort. If (i) $U(c, a) = U(c - g(a))$, (ii) $g'(0) = 0$ and (iii) $\mathcal{A} = \mathbb{R}_+$, the incentive constraint (10) always admits an interior solution for a^* . Hence the recursive condition (14) holds. We will show in the next section that many sequences a_t^* satisfy it. The simplest one is obtained setting a_{t+1}^* equal to a_t^* to derive the rest point solution

$$\hat{a} = g'^{-1} \left(\frac{1 - \lambda}{r^{-1} - \lambda} \right).$$

Efficient effort would require instead that $a = g'^{-1}(1)$. Thus, when r is smaller than one, $\hat{a} < g'^{-1}(1)$ and effort at the rest point is suboptimal because costs are paid upfront, whereas reputational benefits accrue slowly over time. This delay explains why the gap between efficient and actual effort is larger when the agent is more impatient.

¹²Formally, the incentive constraint (11) does not have to hold as an equality at the bounds of the feasibility set. When $\kappa < (1 - \lambda)/(r^{-1} - \lambda)$ and $a_s^* = \bar{a} = \hat{a}$ for all $s > t$, (11) holds as an equality if $a_t^* = \bar{a} + \frac{1}{\gamma(1-\kappa)} \log \left(\frac{\kappa(r^{-1}-\lambda)}{1-\lambda} \right)$ and as an inequality if $a_t^* = \bar{a}$. Hence setting effort equal to $\hat{a}$ in all future periods does not pin-down a unique effort level in the current period. The equilibrium path can converge from below to the rest-point $\hat{a}$ or remain constant at $\hat{a}$. In the latter case, the agent would like to raise her effort but is constrained by the requirement that $a \leq \bar{a}$.

4 Cycles

Having characterized the rest point, we now explain how and when cycles become possible. As in the previous section, we follow a heuristic approach by focusing first on the simpler model with linear costs.

4.1 Cycles with linear costs

Deterministic output.—To build our intuition, we first assume that output is deterministic, i.e., $\sigma_\varepsilon = 0$. Then $\bar{\sigma}_\theta^2 = \sigma_\nu^2$ since all the uncertainty is due to the fact that the agent's type θ changes from period to period. If θ were to remain constant, each type would be revealed once and for all after the first output realization. This simplified environment enables us to isolate the impact of the discount factor. As shown in Proposition 3, the model exhibits stationary cycles of period two when agents are sufficiently patient. Moreover, two-period cycles can Pareto dominate the rest point $\hat{a}$ when it is inefficient, i.e., when $\hat{a} = 0$ because effort costs are too high.¹³

Proposition 3 *Assume that (i) costs are linear in effort ($g(a) = \kappa a$ with $\kappa \in (0, 1)$); (ii) $a \in [0, \bar{a}]$; and (iii) output is deterministic ($\sigma_\varepsilon = 0$). Two-period cycles in which effort oscillates between 0 and $\bar{a}$ become possible if and only if $r \in (\underline{r}, \bar{r})$, where $\underline{r} = \kappa \exp(-\gamma(1 - \kappa)\bar{a})$ and $\bar{r} = \kappa \exp(\gamma(1 - \kappa)\bar{a})$. Moreover, whenever $r \in (\underline{r}, \kappa)$, the rest point $\hat{a} = 0$ is Pareto dominated by cycles.*

Cycles arise because the discount factor fluctuates procyclically. High $a_t = \bar{a}$ entails above-normal consumption, leading to a low marginal utility $U'(c_t - g(a_t))$. The opposite is true next period as $a_{t+1} = 0$ delivers low consumption and relatively high marginal utility $U'(c_{t+1} - g(a_{t+1}))$. Thus $E_t[U'(c_s - g(a_s)) / U'(c_t - g(a_t))]$ is large today, and this justifies the higher investment in reputation building. A similar but opposite mechanism operates in the next period, making low effort optimal. The oscillations of the discount factor capture the willingness of the agent to smooth consumption. Since the benefits of her action only accrue in the following period, she finds it optimal to transfer resources from good to bad times by overinvesting.

¹³Consumers are indifferent between all the different equilibria because they always pay a fair price for the services of the agent. Hence, raising the agent's welfare is equivalent to a Pareto improvement in the overall allocation.

To better understand the forces at work, it is insightful to study the incentive constraint at the upper-bound $\bar{a}$. Full effort is incentive compatible if and only if¹⁴

$$\kappa \leq r \frac{U'(0)}{U'((1-\kappa)\bar{a})} = r \exp(\gamma(1-\kappa)\bar{a}). \quad (17)$$

Reputational returns are on the right-hand side of (17), and they are increasing in effort $\bar{a}$ because it raises the ratio of marginal utilities between bad and good times. Quite intuitively, smaller oscillations become sustainable when the power of incentives is strengthened. Take for example an increase in the discount factor β . It raises r as patient agents tend to be more concerned by their reputation. Hence the required wedge between marginal utilities is decreasing in β which allows the model to generate cycles for lower values of $\bar{a}$. A similar mechanism is triggered by changes in the degree of risk aversion because the effective rate of time preference r is increasing in γ . This effect is reinforced by the positive impact that γ has on the curvature of the marginal utility function. When agents are more prudent, similar oscillations in consumption generate greater swings in the discount factor, which lowers the value of $\bar{a}$ that restores incentive compatibility.

To sum-up, the agent must be sufficiently prudent and not too impatient for cycles to emerge. Note that Proposition 3 also requires that agents are not too patient since it imposes an upper-bound on the effective discount factor ($r < \kappa$). Although cycles can emerge even when this requirement is violated, they will not be welfare improving because, as explained before, the rest point is efficient when costs κ are lower than r .¹⁵

Stochastic output.—We now consider the full model where output is stochastic. When σ_ε differs from 0, today's action affects beliefs on-the-equilibrium path not only in the next period, but also in all future periods. The rate of decay of past observations is controlled by the parameter λ , whose value is equal to 0 when output is noiseless. Proposition 4 takes into account the impact of λ by extending Proposition 3 to environments where output is noisy.

Proposition 4 *Assume that (i) costs are linear ($g(a) = \kappa a$ with $\kappa \in (0, 1)$); and (ii) $a \in [0, \bar{a}]$. There exists an interval $(\underline{r}, \bar{r})$, with $\underline{r} < \frac{\kappa}{\kappa + (1-\lambda)(1-\kappa)}$, such that two-period cycles in which effort oscillates between 0 and $\bar{a}$ become possible if and only if $r \in$*

¹⁴See proof of Proposition 4 for a derivation of (17).

¹⁵Remember that $\lambda = 0$ when $\sigma_\varepsilon = 0$.

$(\underline{r}, \bar{r})$. Moreover, whenever $r \in \left(\underline{r}, \frac{\kappa}{\kappa + (1-\lambda)(1-\kappa)}\right)$, the rest point $\hat{a} = 0$ is Pareto dominated by cycles.

As with deterministic output, welfare improving cycles can arise solely if the agent is sufficiently patient and prudent. We can also show that they require output to not be too noisy. Thus, although σ_ε is, by the usual account, a measure of idiosyncratic uncertainty, higher σ_ε does not necessarily translate into more income volatility for the agent. An increase in σ_ε lowers the power of incentives because recent observations are less informative when realized output is a noisier measure of actual quality. This implies in turn that cycles are not sustainable if σ_ε is too high.

Proposition 5 *There exists a unique output variance $\bar{\sigma}_\varepsilon$ such that cycles of period two are never sustainable when $\sigma_\varepsilon > \bar{\sigma}_\varepsilon$. Moreover, (i) when $\kappa > r$, cycles are sustainable if and only if $\bar{a} \geq \frac{\log(\kappa/r)}{\gamma(1-\kappa)}$ and $\sigma_\varepsilon \in [0, \bar{\sigma}_\varepsilon]$; (ii) when $\kappa \leq r$, there exists an output variance $\underline{\sigma}_\varepsilon \in [0, \bar{\sigma}_\varepsilon)$ such that cycles are sustainable for all $\sigma_\varepsilon \in [\underline{\sigma}_\varepsilon, \bar{\sigma}_\varepsilon]$.*

According to Proposition 5, a reduction in the noisiness of the technology can raise income volatility. This paradoxical behavior occurs if: (i) σ_ε is initially higher than the upper-bound $\bar{\sigma}_\varepsilon$, so that a_t^* is at the rest point; and (ii) the drop in σ_ε is large enough to ensure that the new value is smaller than $\bar{\sigma}_\varepsilon$, so that effort may cycle between 0 and $\bar{a}$. Then the cross-sectional and conditional variance of income will be negatively correlated with σ_ε . As mentioned in the introduction, this relates to the evidence on rising inequality in SPGBW.

The variance of income increments from period to period is given by¹⁶

$$\begin{aligned} \text{Var}(m_{t+1} + a_{t+1}^* - (m_t + a_t^*)) &= (1 - \lambda)^2 (\sigma_\varepsilon^2 + \bar{\sigma}_\theta^2) + \text{Var}(a_t^*) \\ &= \sigma_\nu^2 + \text{Var}(a_t^*). \end{aligned} \quad (18)$$

The variance of a_t^* is obviously equal to 0 at the rest point. It can jump to $\bar{a}^2/4$ when σ_ε crosses the $\bar{\sigma}_\varepsilon$ threshold from above if the agent selects the cycling equilibrium. The drop in output uncertainty could therefore trigger an increase in the conditional volatility of income.

To characterize its impact on the cross-sectional variance, we must first ensure that a stationary distribution of types exists. The simplest way to do so is to assume that agents randomly exit the market with probability δ per period. Since each agent's

¹⁶See the proof of Proposition 2 in the Appendix for a derivation of the second equality.

Figure 1: CROSS-SECTIONAL EARNINGS VARIANCE AS A FUNCTION OF σ_ε .

Notes: $\gamma = 3, r = .8, \kappa = .9, \sigma_v = .4, \delta = .3, \mathcal{A} = [0, 1]$.

productivity is a geometrically distributed random variable and an additional period of life adds σ_v^2 to the variance of m , the *time-invariant* distribution of types $\Upsilon(m)$ is a mixture of normal distributions with variance: $\text{Var}(m_t) = \sigma_v^2 \sum_{t=1}^{\infty} t \delta (1 - \delta)^t$.¹⁷ Given that income variance is equal to $\text{Var}(m_t) + \text{Var}(a_t^*)$, raising σ_ε above the $\bar{\sigma}_\varepsilon$ threshold can reduce income variance. Provided that agents were coordinated on the reputation cycle, an increase in σ_ε will trigger a switch to the rest point and a fall in the cross-sectional variance, as depicted in Fig. 1. The impact of the other volatility parameter σ_v is less intriguing because it strengthens the power of incentives. It is therefore positively correlated with income variance since a switch to the cycling regime will be triggered by an increase in σ_v , as opposed to a decrease in σ_ε . Note that σ_v also directly increases the volatility of income by augmenting the variance of income increments from period to period (see Ed. (18)). As a result σ_v raises the discount factor r , which further supports the emergence of cycles.

4.2 Cycles with Convex Costs

We now turn our attention to the model with convex costs introduced in Section 3.2. We first analyze the stability conditions under which all paths starting in the vicinity

¹⁷More precisely, we know from (18) that $m_{t+1} \sim \mathcal{N}(m_t, \sigma_v^2)$. Hence, the cross-sectional variance of m among the cohort of agents that have been active for t periods is equal to $t\sigma_v^2$. The overall variance of m is obtained summing up the variances across all cohorts and weighting them by the fraction, $\delta(1 - \delta)^t$, of agents in each cohort t .

of the rest-point converge back to it. Then we show that, when the cost function is quadratic, deterministic 2-period cycles are sustainable whenever the rest point is backward unstable.

Stability of the rest point.—Since today's action a_t^* is determined by tomorrow's action a_{t+1}^* , we analyze the backward dynamics of the model. Hence we say that the rest point is locally unstable when the paths that land in its vicinity gradually diverge from the rest point as time unfolds backwards. Under risk neutrality, the rest point is always backward stable.¹⁸ Any other anticipations than $a_{t+1}^* = \hat{a}$ generate diverging trajectories, hence the unique deterministic equilibrium requires that $a_t^* = \hat{a}$ for all t . By contrast, the rest point can be backward unstable when the agent is risk averse. Then, for any initial anticipation a_1 in the neighborhood of $\hat{a}$, there exists a deterministic path that converges to $\hat{a}$. The condition under which equilibrium multiplicity arises is laid-out in Proposition 6.

Proposition 6 *Assume that $g(a)$ is strictly convex with $g'(0) = 0$. Then the rest point $\hat{a}$ is backward unstable if and only if*

$$\left(\frac{1 - \lambda + 2\lambda g'(\hat{a})}{1 - \lambda + \lambda g'(\hat{a})} \right) \frac{g''(\hat{a})}{g'(\hat{a})} < 2\gamma(1 - g'(\hat{a})). \quad (19)$$

Equation (19) depends on the action $\hat{a}$ at the rest point, and so only implicitly describes the parameter values under which multiple equilibria arise. We can nonetheless easily establish a couple of observations. First, when the coefficient of absolute risk aversion γ goes to zero, condition (19) is violated and the rest point is backward stable.¹⁹ Second, if $\hat{a}$ is efficient, the rest point is again backward stable. To see why, note that the impact of γ , as measured by the term on the right hand side of (19), is proportional to $1 - g'(\hat{a})$. Thus risk aversion is irrelevant when $g'(\hat{a}) = 1$, that is when the action is efficient. Intuitively, the curvature of the utility function matters only to the extent that changes in a affect utility. Benefits and costs are set equal at the first best and a marginal increase in effort leaves utility unchanged. However, we have seen that $\hat{a}$ is inefficient whenever the agent discounts the future. Then raising

¹⁸Formally, let $\varphi(a_{t+1}^*) = a_t^*$ denote the implicit map resulting from (14). The rest point is stable when the absolute value of the derivative of $\varphi(a_{t+1})$ at the rest-point is less than one, i.e., when $|d\varphi(\hat{a})/da| < 1$. See Fig. 2 for an illustration of the model's dynamics under risk neutrality.

¹⁹This is not obvious from (19) because γ affects $\hat{a}$ through its impact on r . However, since r converges to β when γ goes to zero, $\lim_{\gamma \rightarrow 0} \hat{a}(\gamma) = g'^{-1}((1 - \lambda) / (\beta^{-1} - \lambda))$. Hence the LHS of (19) converges to a positive value, whereas the RHS goes to zero, and so (19) never holds in the limit.

a lowers her marginal utility, thus making it attractive to provide more effort today in order to raise tomorrow's consumption. This is why an increase in effort is followed by a drop and a return to the rest point if the curvature of the utility function is high enough.

Cycles.—When the rest point is backward unstable, there exists a continuum of deterministic equilibria whose paths start in the vicinity of the stationary solution and gradually converge to it. But the model can also generate more complex behaviors with persistent oscillations. We now focus on the simplest equilibria in such class, namely deterministic cycles of period two.

Proposition 7 *When the costs function is quadratic, i.e., $g(a) = a^2/2$, deterministic 2-period cycles are sustainable whenever the rest point is locally unstable.*

The reason why cycles are sustainable is the same as in the model with linear costs: Procyclical fluctuations in the discount factor make it optimal to work hard when consumption is high and to reduce effort when consumption is low. This mechanism is closely related to the one that renders the rest point backward unstable, and Proposition 7 shows that the two phenomena are indeed concomitant when disutility is quadratic in effort.

Inspecting the phase portrait generated by the incentive constraint (13) helps one understand how cycles arise. We highlight the impact of risk aversion by also reporting the model's dynamics when agents are risk neutral. Then, as shown in the left-hand panel of Fig. 2, the mapping between a_t^* and a_{t+1}^* is linear with a slope greater than one. Hence the rest point $\hat{a}$ is the only effort level that does not generate diverging action paths.

We introduce risk aversion in the right-hand panel of Fig. 2. The dynamic map solves the recursive equation (14) with a quadratic cost function $g(a) = a^2/2$, and is therefore equivalent to

$$\exp(-\gamma(a_t^* - a_t^{*2}/2)) a_t^* = r \exp(-\gamma(a_{t+1}^* - a_{t+1}^{*2}/2)) [1 - \lambda + \lambda a_{t+1}^*]. \quad (20)$$

Let us first explain why small values of a_t^* can never be incentive compatible. Letting a_t^* go to zero, we see that the expression on the left hand side of (20) also converges to zero. By contrast, the expression on the right hand side of (20) has a positive minimum. By continuity, there always exists an $\underline{a} > 0$ such that (20) does not

Figure 2: PHASE PORTRAITS WITH LINEAR AND CARA UTILITY.

Notes: $\lambda = .7$, $\beta = .85$, $\gamma = 5$, $g(a) = a^2/2$.

hold whenever $a_t^* \in [0, \underline{a}]$. In economic terms, there is no expectation about a_{t+1}^* that sustains an effort level smaller than $\underline{a}$. Intuitively, the dynamic map in Fig. 2 is empty close to the origin because today's marginal loss converges to 0 whereas tomorrow's marginal returns are bounded from below by $1 - \lambda$.

Besides this empty interval, we see that, instead of the one-to-one mapping prevailing under risk neutrality, the dynamic map under risk aversion is a correspondence which associates a pair of incentive compatible a_{t+1}^* to any a_t^* . The incentive constraint is satisfied by two different a_{t+1}^* because tomorrow's effort shifts the discount factor and the marginal costs in opposite directions. Increasing a_{t+1}^* raises tomorrow's consumption, which lowers the discount factor and counteracts the increase in marginal costs $g'(a_{t+1}^*)$. When the elasticity of the discount factor with respect to a_{t+1}^* is higher than that of the marginal costs, it is possible to perturb a pair of sustainable actions and restore incentive compatibility by adjusting a_{t+1}^* until its effect on the discount factor offsets the change in marginal costs.

Effort settles down at the rest point $\hat{a}$ when the expectation-formation mechanism always selects the lower branch of the phase portrait. By contrast, 2-period cycles are *regime switching* equilibria whose actions oscillate between the upper and lower branch.

Figure 3: BACKWARD INSTABILITY OF THE REST POINT $\hat{a}$ AS A FUNCTION OF γ , β AND λ .

Stability region.—Equation (19) enables us to identify the region with multiple equilibria. The rest point is backward unstable and 2-period cycles arise whenever the parameters fall inside the surface depicted in Fig. 3. We use two different angles to depict the instability region because its shape is strongly asymmetric.

Looking first at the discount factor β , we see that it has an ambiguous impact as cycles emerge only for intermediary values of β . When β goes to 0, so does the rest point $\hat{a}$ because the agent is not anymore concerned by her reputation. Hence the marginal disutility of effort $g'(\hat{a})$ also converges to 0, and the instability condition (19) is always violated. Reputational returns are so low that an increase in today's effort can be incentive compatible solely if the marginal disutility of effort in the next period increases more than proportionally. Thus the slope of the relationship between a_{t+1}^* and a_t^* is smaller than 1 when evaluated at $\hat{a}$, which ensures that the rest point is locally stable. Conversely, as β goes to 1, $\hat{a}$ converges to the efficient level of effort which, as explained above, is always backward stable.

Comparing the impact of the coefficient of risk aversion γ in the right and left-hand panels of Fig. 3, we notice that γ has a stronger effect on the instability region when β is low. This is because γ does not only increase risk aversion but also renders the agent more prudent. As explained before, the precautionary motive increases the

level of effort at the rest point. When $\hat{a}$ is already close to its efficient level because β is high, as in the right-hand panel of Fig. 3, then the precautionary channel counteracts the risk aversion channel. By contrast, when β is low, as in the left-hand panel of Fig. 3, the risk aversion channel dwarfs the precautionary channel, so that raising γ makes it much more likely for cycles to be sustainable.

The last variable included in Fig. 3 is the gain parameter λ which controls the inertia of beliefs. Remember that varying λ is equivalent to adjusting the output variance σ_ε^2 , with σ_ε^2 going from 0 to infinity as λ goes from 0 to 1.²⁰ As already discussed in Section 4, decreasing λ makes it more likely that cycles arise by strengthening the power of incentives. However, it also increases the value of the rest point, which may counteract the direct impact of λ when the rest point is close to its efficient level. As was the case for γ , this countervailing force explains why the instability region is much less sensitive to changes in λ in the right-hand panel of Fig. 3, where β and thus $\hat{a}$ are both close to one.

Output volatility and income variance.—Studying the model with linear costs, we proved that a reduction in the noisiness of the technology may raise income variance. A similar conclusion can be numerically established in the model with convex costs. For instance, taking the parameters reported below in Table 1 and varying the value of σ_ε , we indeed find that income variance and output volatility can be negatively correlated when effort fluctuates endogenously.

The upper panel of Fig. 4 reports effort in the high and low states of 2-period cycles along with its value at the rest point. Cycles are sustainable when incentives are strong enough. Hence they arise when λ , and consequently σ_ε , are below a certain threshold.²¹ The cross-sectional variance of income increases as σ_ε crosses this bifurcation point because effort is not anymore the same across all agents. The bottom panel is the analog of Fig. 1; this experiment confirms that an increase in income variance can be caused by a *reduction* in the volatility of the underlying shocks. Moreover, we see that the calibrated value of σ_ε in Table 1 falls well within the region where cycles are sustainable. This again relates to the evidence on rising inequality in SPGBW.

²⁰By contrast, the variance of types σ_ν^2 affects two parameters at once. Raising σ_ν^2 leads to a decrease in λ as well as an increase in r , which is equivalent to raising β in Fig. 3. In general, these two adjustments tend to reinforce each other, thus making it more likely that cycles are sustainable.

²¹Note that Fig. 3 shows that it is not always possible to generate cycles by lowering λ , especially when β is close to one so that the rest point is not too far from its efficient level.

Figure 4: EQUILIBRIUM EFFORT AND EARNINGS VARIANCE AS A FUNCTION OF σ_ε .
 Notes: Other parameters reported in Table 1.

5 Earnings Persistence

5.1 Deterministic cycles

Parametrization.—We have established that reputation cycles become possible when the risk-aversion and the patience of the agent belong to some intermediate range. To see whether these two requirements are likely to be met in practice, we use micro data to pin down the parameters.

We now assume that agents are workers who maximize their expected utility, and that a is a worker's hidden labor supply. We focus on quadratic disutility so that $g(a) = a^2/2$ so that the model has only five parameters $\{\gamma, \beta, \delta, \sigma_\varepsilon, \sigma_\nu\}$. For the yearly discount factor, we set $\beta = 0.95$. The exit rate δ determines the average length of a worker's career. Assuming that reputation is occupation specific, we target Kambourov and Manovskii's (2008) finding that occupational mobility at the one-digit level averaged 13% per year in PSID data over the 1968-1997 period.

The two variance parameters $\{\sigma_\varepsilon, \sigma_\nu\}$ are chosen to match the speed of employer learning and the labor wedge. We rely on the labor wedge because it is the most widely documented measure of labor market inefficiency. Let τ denote the labor wedge which we define as the gap between the marginal product of labor (MPN) and the marginal rate of substitution (MRS), so that $\tau \equiv 1 - \text{MRS}/\text{MPN}$. The labor wedge measures the extent to which agents underinvest in their reputation. In the absence of

distortions, $\tau = 0$ because both MRS and MPN are equal to the wage. At the rest point of our model, however, labor supply is inefficiently low and the MRS is lower than the wage. Hence the labor wedge is positive and equal to²²

$$\hat{\tau} = \frac{r^{-1} - 1}{r^{-1} - \lambda} \in (0, 1), \quad (21)$$

where λ and r are defined in (7) and (12). We write $\hat{\tau}$ for the wedge when labor supply is at the rest point so that $a = \hat{a}$. We set $\hat{\tau} = 0.4$ to match the average labor wedge in the US over the last fifty years, as documented by Shimer (2009).

For our second moment, we target the frequency at which the market updates its beliefs, i.e., the speed with which information about the agent's performance leaks out to the general public. Its value has been estimated by Lange (2007) in his empirical analysis about the speed of employer learning. Using AFQT scores to infer ability, Lange's estimates of the learning speed implies that $\lambda = .71$.²³

Together with Eqs. (7), (21), these estimates for λ and τ enable us to pin down σ_ε and σ_ν conditional on γ . The upper panel of Fig. 5 shows that reputation cycles are consistent with our calibration strategy whenever $\gamma \geq 3.3$. The amplitude of the oscillations increases with γ , as greater fluctuations are required to offset the increase in the curvature of the utility function. The lower panel of Fig. 5 shows that effort at the rest point is always smaller than average effort in the cycling regime, thus substantiating our claim that cycles can incentivize investment in reputation.²⁴ Although γ is a CARA-utility parameter, we can convert the requirement that $\gamma \geq 3.3$ into a statement about CRRA utility as follows: By adjusting average ability $\bar{\theta}$, we can normalize average consumption and output to unity at the rest point.²⁵ Then γ also corresponds to

²²The wage can be normalized to one without loss of generality. The technology being linear in supplied labor a , the MPN is also equal to one and so the labor wedge at the rest point is equal to $\hat{\tau} = 1 - \text{MRS}/\text{MPN} = 1 - g'(\hat{a}) = 1 - \frac{1-\lambda}{r^{-1}-\lambda}$, i.e., (21). Note that the macro and micro labor wedges are identical in steady-state because all agents behave in the same way. Hence our definition is consistent with the aggregate measure provided by Shimer (2009).

²³Lange (2007) uses K_1 to denote the speed of learning parameter and defines it on page 11, eq. (11). Comparing its expression with that of λ in (7), one sees that $K_1 = 1 - \lambda$. Although learning in Lange's model is non-stationary, as types remain constant over time, the similarity between K_1 and λ holds because K_1 captures the speed of learning during the first year of a worker's career. Since ability is occupation specific in our calibration, we have to correct the empirical estimate $\hat{K}_1$ by the rate at which workers switch occupation, as $\hat{K}_1 = K_1(1 - \delta) = (1 - \lambda)(1 - \delta)$.

²⁴Note that cycles are not always welfare improving. Alternative calibrations show that they can also push average effort below its value at the rest point.

²⁵More precisely, setting $\bar{\theta} = 0.4$ ensures that consumption at the rest point is on average equal to 1 since $E[\hat{c}] = \bar{\theta} + \hat{a} = \bar{\theta} + 1 - \hat{\tau}$, and $\hat{\tau}$ has been set equal to 0.4 to match the evidence in Shimer (2009).

Figure 5: DETERMINISTIC CYCLES AS A FUNCTION OF THE COEFFICIENT OF RISK AVERSION γ .

Notes: Baseline parameters reported in Table 1.

the average coefficient of relative risk aversion. Picking $\gamma = 4$, while setting $\sigma_\varepsilon = .117$ and $\sigma_\nu = .04$, ensures that we exactly match our labor wedge and speed of learning targets. Our choice of parameters is summarized in Table 1.

Earnings persistence.—We now explain why reputation cycles generate earnings profiles with large shocks that are less persistent than small shocks, as documented in GKOS and ABB. To relate mean reversion to the size of the shocks, we follow GKOS in reporting today's earnings changes against the expected value of future earnings changes k years from now, for different values of k . The resulting impulse response functions are depicted in Fig. 6. They measure the fraction of each shock that has mean-reverted. The points that are on the horizontal line indicate perfect persistence of today's shocks, whereas those that are on the 45 degrees line indicate full mean reversion.²⁶

We first report in the left-hand panel of Fig. 6 the impulse responses when effort is at the rest point. Then earnings shocks, whatever their sizes, are fully persistent because beliefs, and thus wages, follow a martingale process. We introduce reputation cycles in the right-hand panel of Fig. 6. Large shocks are now driven by changes in

²⁶The impulse response functions are estimated using a simulated panel dataset made of 100 waves, each one of them containing 5.000 workers. We delete the first 50 waves to limit the impact of the initial conditions, thus basing our estimates on 500.000 observations.

Table 1: Calibrated Parameters

DIRECTLY CALIBRATED PARAMETERS					
PARAMETER [†]	VALUE		INTERPRETATION	SOURCE	
β	.95		Discount factor	Standard	
θ	.4		Mean initial ability	Normalize consumption $\hat{c} = 1$	
δ	.13		Career switching rate	Kambourov et al. (2008)	
INDIRECTLY CALIBRATED PARAMETERS					
PARAMETER [†]	$\gamma=4$	$\gamma=9.1$	INTERPRETATION	TARGET	SOURCE
σ_ε	.117	.051	Output noise	Labor wedge $\hat{\tau}= .4$	Shimer (2009)
σ_ν	.040	.017	Type volatility	Speed of learning $\lambda = .71$	Lange (2007)
INDIRECTLY CALIBRATED PARAMETERS OF STOCHASTIC MODEL					
PARAMETER [†]	VALUE		INTERPRETATION	SOURCE	
σ_ξ	.124		Volatility observable productivity	Impulse response of log earnings changes	Guvenen et al. (2015)
α	.75		Switching probability	Impulse response of log earnings changes	Guvenen et al. (2015)
γ	9.1		Coefficient of risk aversion	Impulse response of log earnings changes	Guvenen et al. (2015)

Notes. [†]When applicable, parameter values are for yearly frequency.

effort, with cycles considerably widening the range of income fluctuations. Focusing on horizons that are even ($k = 1, 3, 5$), large shocks are completely transitory because they follow the agent's effort whose value oscillates from one period to the next.

5.2 Three fit-improving extensions

The following three extensions improve the model's fit to the impulse-response evidence in GKOS.

1. **Observed characteristics:** Given that some output characteristics may be observable, we add a state p_t such that the output of a worker with t years of occupational experience reads

$$y_t = \underbrace{\theta_t + a_t + \varepsilon_t}_{Unobservable} + \underbrace{p_t}_{Observable}, \quad (22)$$

where

$$p_t = p_{t-1} + \xi_t, \text{ with } \xi_t \sim \mathcal{N}(0, \sigma_\xi^2) \text{ and } p_1 = 0.$$

Since p_t is persistent, it drops out of the incentive constraints for the same reason that m_t does.

2. **Some workers can borrow and lend:** Cycling equilibria arise only if agents cannot borrow and lend. According to Kaplan, Violante and Weidner (2014) only

Figure 6: IMPULSE RESPONSES OF DETERMINISTIC CYCLES MODEL.

Notes: Using the parameters in Table 1 when $\gamma = 4$, we generate a balanced panel dataset of 50 waves, with 5,000 workers per wave. For each period, we assign workers into 10 groups based on their current log-earnings growth. Then, for workers in each group, we compute the average of log k -year earnings growth, i.e., $E[\log(c_{t+k}/c_t) | \log(c_t/c_{t-1})]$ for $k = 1, 2, 3, 4, 5$.

one third of agents are so constrained. Hence we assume that two thirds of workers can smooth their income and implement the rest point solution since that is the only deterministic equilibrium that exists in that case. This adds no new parameters.

3. **Stochastic cycles:** Our model exhibits a significant departure from GKOS data as the degree of mean reversion dramatically changes when the forward horizon k goes from even to odd. Whenever k is odd, shocks are fully persistent because effort reverts to its original value. This suggests that, in order to reproduce the data, one needs to consider cycles whose lengths are not deterministic. Thus we allow agents to switch stochastically, instead of deterministically, between periods of low and high effort. Signals that raise a_t are “good news” because they bring the worker closer to first best, signals that lower a_t are bad news.²⁷ We continue to focus on action paths with two levels only, but assume that, at each level, effort changes state with probability α , or remains constant with the complementary probability. This lowers the discount factor but other-

²⁷This relates to the good-news and bad-news literature. By contrast to our signals, Board and Meyer-ter-Vehn (2013), e.g., have endogenous signals to which effort can respond.

wise leaves incentives as before; details are in Appendix A.4. One coordination device may be task assignment as in Bar-Isaac and Hörner (2014) – assignment to a demanding task indicating an expectation that effort will be high, and vice versa.

5.3 Extended calibration

We keep the baseline parameters reported in Table 1, and select the values of the coefficient of risk aversion γ and of the two new parameters $\{\alpha, \sigma_\xi\}$ that minimize the distance between the persistence of earning shocks predicted by the model and their empirical counterparts in GKOS.²⁸ As reported in the third column of Table 1, we obtain the best fit when $\alpha = .75$, $\sigma_\xi = .124$ and $\gamma = 9.1$.

Distribution of earnings growth.—The density of earnings changes predicted by the model, and the one observed in GKOS data, are reported in the upper-panel of Fig. 7. We also include a normal density with the same standard deviation as in the data. Fig. 7 shows that the simulated and empirical densities display high kurtosis since they have much sharper peaks than the normal distribution. In our model, the extra kurtosis is generated by reputation cycles because the two other sources of income shocks, namely changes in reputation and observable productivity, are independently drawn from normal distributions. The impact of reputation cycles is noticeable on the shoulders of the simulated distribution, that is in the regions around the positive and negative standard deviations, where the model generates slightly more mass than in the data.

Persistence of earnings changes.—The impulse responses of median income workers are reported in Fig. 7. The prediction of the model are now very much in line with GKOS data. First, the impulse responses display similar shapes. According to our model, the nonlinearity of the impulse response is explained by the fact that small to medium shocks are very persistent because they result from changes in reputation or observable productivity, which are permanent by assumption. By contrast, large shocks are driven by reputation cycles that tend to revert over time. Given

²⁸More precisely, the empirical moments are summarized by the vector $\hat{\mathbf{m}}$ which contains the average persistence of earnings shocks at an horizon of 3 years for 20 equally-sized quantiles. Using $\mathbf{m}(\mathbf{v})$ to denote the simulated moments generated by the vector of parameters $\mathbf{v} \equiv \{\alpha, \sigma_\xi, \gamma\}$, we compute the quadratic distances $d(\mathbf{v}) \equiv (\mathbf{m}(\mathbf{v}) - \hat{\mathbf{m}}) \cdot (\mathbf{m}(\mathbf{v}) - \hat{\mathbf{m}})$. Our calibrated vector $\hat{\mathbf{v}}$ solves $\hat{\mathbf{v}} = \arg \min_{\mathbf{v} \in [0,1] \times \mathbb{R}_+^2} d(\mathbf{v})$.

Figure 7: HISTOGRAM OF LOG EARNINGS CHANGES AND IMPULSE RESPONSES OF STOCHASTIC CYCLES MODEL.

Notes: Parameters reported in Table 1. Following GKOS, we focus on the impulse response of median income workers, i.e. workers whose income at $t - 1$ falls between the 46th and 55th percentile.

that the empirical relevance of large transitory shocks could already be inferred from GKOS, our contribution consists in providing a microfoundation for their emergence. Adding a stochastic switching rate ($\alpha = .75$) smooths the impulse responses, as one-fourth of saving-constrained agents do not adjust their effort from one period to the next.

Besides fitting the persistence of earnings shocks at each quantile, our simulations also approximate the size of the shocks, $\log(c_t) - \log(c_{t-1})$, reported on the horizontal axis of the lower panel of Fig. 7. Achieving a good fit along both dimensions was made possible by the adjustment of the coefficient of risk aversion γ . As shown in Fig. 5, reputation cycles are wider when γ increases. Our calibration indicates that matching the actual size of earning shocks requires setting the coefficient of risk aversion γ around 9.

Earnings volatility.—According to our model, hand-to-mouth agents should exhibit higher earnings volatility since they cycle between periods of low and high effort. This prediction is broadly consistent with the empirical findings in Ziliak et al. (2011) and Dahl et al. (2011) according to which earnings volatility is much higher among high school dropouts. They find that differences in earnings volatility across educational groups are mostly accounted for by employment transitions, and argue that further research should strive to assess whether the employment transitions that lead to higher volatility are voluntary or not. Reputation cycles support the first interpretation since high school dropouts are more likely to be saving-constrained, and as such may find it optimal to supply different amounts of labor over time.

6 Discussion

6.1 Related models

We have shown that reputation cycles capture the non-linear persistence of income shocks. To the best of our knowledge, no other explanation has yet been advanced. Standard models of income dynamics may account for the non-linear persistence of either positive or negative shocks, but not both as our model does. Take for instance search models with on-the-job search. The loss of a high-paying job can trigger a large income loss that is not fully persistent as workers can search their way out of

low paying jobs. A positive income shock, on the other hand, is more persistent since it is triggered by the find of a higher-paying job as in Burdett and Mortensen (1996) for example. Hence on-the-job search does not explain the transitory nature of large income gains. A natural candidate would be performance based bonuses since exceptional rewards are not likely to be awarded over two consecutive years. However, bonuses being added to baseline earnings, they cannot account for the low persistence of large negative shocks.

Our analysis demonstrates that oscillating effort provides a unified explanation for both negative and positive shocks. Hence one may wonder whether canonical models of human capital accumulation are able to generate effort paths that are observationally equivalent to reputation cycles. We now explain why this is not the case by isolating the two unique features of our model:

(i) effort is below its output-maximizing level so that higher effort today *raises* current earnings and utility,

(ii) higher effort today does *not* raise future output.

Properties (i) and (ii) combine to generate the movements in the effective discount factor that are needed to sustain cyclical effort paths. By contrast, models of human capital investment do not feature oscillating solutions because they include at most one, but not both of these features.

Ben-Porath type of human capital investment.—Reputation, as we have modeled it, differs from traditional models of the Ben-Porath (1967) type in which a worker sacrifices current income in return for higher future earnings. In contrast to (i) and (ii) such human capital investment lowers current earnings via a compensating differential contrary to property (i), and it raises future earnings contrary to property (ii).

Learning by doing investment.—Although it does not have property (ii), LBD is the closest-related line of research because it potentially shares with our model property (i), namely the feature that higher learning today raises an agent's income today. That does not mean, however, that it also raises today's utility. Here is a version of LBD that compares closely to our model. Let k denote the stock of human capital

and a the investment in human capital. Assume that

$$c_t = k_t + a_t \text{ and,} \quad (23)$$

$$k_t = (1 - \delta)k_{t-1} + a_{t-1}, \text{ with } \delta \in (0, 1). \quad (24)$$

Property (i) does not hold because agents set a_t beyond the point where today's utility $U(c_t - g(a_t))$ is maximized.²⁹ Hence increasing a_t *reduces* current utility instead of raising it. Property (ii) does not hold either because, according to (24), higher effort raises k' , and so increases tomorrow's output. Thus the discount factor is decreasing in a_t at the rest point, and LBD exerts a stabilizing force which prevents the emergence of cycles.

When combined with technological choice as in Parente (1994) and Jovanovic and Nyarko (1996), the LBD model can generate cycles resembling the asymmetric cycles displayed in Fig. 9. It only generates large transitory negative shocks to earnings growth, but not the corresponding positive ones that the evidence in Fig. 7 shows.

6.2 Extensions

1. *Non-stationary learning and life-cycle profiles.*—We now describe four extensions. As a first check, we let the precision of the starting prior $\sigma_{\theta,0}^{-2}$ differ from its long-run level $\bar{\sigma}_{\theta}^{-2}$. We typically expect the precision of beliefs to increase over time, as is always the case when types are constant since then σ_{ν}^{-2} , and thus $\bar{\sigma}_{\theta}^{-2}$, are both infinite. The returns to effort are decreasing in posterior precision because it lowers the weight that the market puts on recent observations. Hence the action path is also decreasing over time as in Holmström's model with risk-neutral agents. The negative correlation between average earnings and labor market experience explains why we did not use our model to analyze life-cycle profiles but instead focused on earnings changes.

Adding risk-aversion allows us to construct another equilibrium sequence which oscillates around the regularly decreasing path, as shown in Fig. 8. This alternative solution eventually coincides with the two-period cycles described in the previous

²⁹Current output is maximized when $a = g^{-1}(1)$. Thus the observation that effort exceeds its output-maximizing level directly follows from the first-order condition

$$g'(a_t) = 1 + \sum_{s=t+1}^{\infty} \beta^{s-t} \frac{U'(c_s - g(a_s))}{U'(c_t - g(a_t))} \frac{\partial k_s}{\partial a_t} > 1. \quad (25)$$

Figure 8: NON-STATIONARY LEARNING GENERATING EVEN WIDER TAILS AT YOUNGER AGES THROUGH CYCLING EFFORT.

Notes: $g(a) = a^2/2$, $\beta = .75$, $\gamma = 3$, $\sigma_\varepsilon = .1$, $\sigma_\nu = .05$, $\sigma_{\theta,0} = .18$.

section because the precision converges to its stationary level $\bar{\sigma}_\theta^{-2}$. Moreover, the decreasing amplitude of the oscillations is consistent with empirical evidence about the negative relationship between income volatility and worker age. Even along the regularly decreasing “Holmström” path shown by the solid curve in Fig. 8, income volatility is decreasing with the worker age. Cycling a_t paths thus reinforce this decline by subtracting mass from the tails of the growth rates’ distributions as the worker ages.

2. *Asymmetric cycles and income-growth persistence.*—Although we have focused on cycles of period 2, the model can generate more complex patterns. For example, it is possible to use the model with linear costs to generate cycles with more than two states. As an illustration, Fig. 9 describes incentive-compatible cycles of period 3 where investment grows for two successive periods and then drops back to its initial level. This is the simplest instance of asymmetric cycles featuring protracted booms and sudden busts. A tendency for a time series to show large negative growth rates followed by several smaller positive growth rates is known as “steep asymmetry.” Moreover, the long-run frequency distribution of a_t places equal weight on the points 0, 0.71, and 1, and thus has longer left tail. A tendency for a detrended time series to have negative skew is known as “deep”.³⁰

The longer period of growth is sustainable because investment costs κ are low, i.e.,

³⁰See Fig. 4 in Jovanovic (2006) for a graphical distinction between the two types of asymmetry.

Figure 9: ASYMMETRIC CYCLES WITH CARA UTILITY AND LINEAR COSTS.

Notes: $\sigma_\varepsilon=.1$, $\sigma_\nu=.2$, $\beta=.8$, $\gamma=4$, $\kappa=.75$, $\mathcal{A}=[0,1]$.

$\kappa < (1 - \lambda)/(r^{-1} - \lambda)$.³¹ As explained in Section 4, when this inequality holds, the rest point is efficient because the agent always finds it profitable to raise her investment. But now she is discouraged from doing so in the two low states by the fact that the marginal utility of consumption will decrease in the following period. We show in Appendix A.3 that this feature gives rise to asymmetric cycles similar to the one depicted in Fig. 9 whenever: (i) $\bar{a}$ is high enough, and (ii) $\kappa \in (\underline{\kappa}, (1 - \lambda)/(r^{-1} - \lambda))$ for some $\underline{\kappa} > 0$. The logic of the proof can be extended to cycles with more than 3 states.

3. *Finite horizon does not eliminate cycles.*—In contrasting our model to Holmström’s benchmark, we have stressed that equilibrium uniqueness is not anymore guaranteed. But one may argue that this distinction is not fundamental because it stems from our focus on deterministic equilibria. Although this requirement is implicitly assumed in much, if not all, of the career concerns literature, there can be other solutions if one allows for strategies that are history-dependent. However, whenever the agent has a finite horizon, backward induction bites and the equilibrium under risk neutrality is unique and deterministic. By contrast, multiple equilibria and reputation cycles continue to arise in finite-horizon models with risk averse agents.³²

³¹By contrast, asymmetric cycles with protracted slumps and sudden booms can be constructed when the steady-state is inefficient, so that $\kappa > (1 - \lambda)/(r^{-1} - \lambda)$.

³²This is particularly simple to check when disutility is linear. Let T denote the agent’s lifetime and assume that, as in Proposition 4, $\kappa \in (\frac{1-\lambda}{r^{-1}-\lambda}, 1)$. Then $a_t^* = 0$ for all $t \in \{1, 2, \dots, T\}$ is a deterministic equilibrium. But it is not always unique. For example, $a_1^* = \bar{a}$ and $a_t^* = 0$ for all $t \in \{2, \dots, T\}$ is also a deterministic equilibrium if the horizon T and the upper bound $\bar{a}$ of the action set are large enough.

4. *CRRA utility function*.—We have restricted our attention to CARA utility functions because they do not display any wealth effect. As shown in the proof of Proposition 1, this property implies that optimal effort does not depend on the agent’s reputation, which enables us to demonstrate that the necessary condition (11) is also sufficient. Establishing sufficiency is notoriously complicated in models with persistent private information such as ours. Given that any deviation drives a long-lasting wedge between the agent’s and the market’s beliefs, a strategy that is not attractive on the equilibrium path might nonetheless be profitable off path. In other words, ruling out one-shot deviations is not sufficient since it leaves open the possibility that multiple deviations are profitable. Cisternas (2018) addresses this issue by deriving necessary and sufficient conditions that hold across a broad class of nonlinear utility functions. Using these conditions, Cisternas establishes the existence of equilibria where the agent’s equilibrium actions only depend on the market’s belief. A promising direction for further research would be to combine Cisternas’ approach with ours, and assess the existence of equilibria where the agent’s actions depend not only on the market’s belief but also, as in our model, on an exogenous signal.

The CARA restriction has therefore been imposed for technical reasons and is not instrumental in generating endogenous fluctuations. If anything, introducing a wealth effect would render agents more willing to take risks when their action and consumption are high, thus reinforcing self-sustaining fluctuations in effort. This was illustrated in the previous version of our paper (Jovanovic and Prat, 2016) where the model was embedded into a general equilibrium setting. There we explained how one recovers the same incentive structure if firms operate the technology described in (1) and if risk-averse households fully diversify their purchases of goods over firms. There are two fundamental differences between the partial and general equilibrium settings. First, in the latter, cycles are in *aggregate* consumption. Hence the requirement that agents cannot smooth them becomes equivalent to the assumption that the economy is closed with no access to outside finance. Second, agents being member of a large representative household, they are insured against idiosyncratic shocks to their wealth, which neutralizes the wealth effect for utility functions outside of the CARA family. We show in Jovanovic and Prat (2016) that a calibrated version of the macro model with CRRA utility produces realistic fluctuations in terms of peak-to-trough movements in consumption and of the spacing of time between recessions, thus illustrating that reputation cycles do not hinge on the CARA specification.

7 Conclusion

We have shown that investment in reputation is not always discouraged by financial constraints. Cyclical paths are sustainable only if the agent cannot borrow or lend, and they may entail higher earnings because average effort can be higher. This contrasts sharply with the findings reported in a host of papers that financial constraints inhibit tangible investment. It is therefore possible, that when facing such constraints, agents may substitute away from investing in physical capital and towards reputation capital. The conditions under which cycles can emerge are stated in Propositions 3 and 6, and the calibrated parameter values satisfy them comfortably.

We have also shown that cyclical effort paths generate unusual shocks to earnings and can help explain the evidence on their dynamics. By inducing a large switch in effort, tail events in earnings growth emerge. As documented in the data, these large shocks are less persistent than the smaller shocks to earnings growth. Given the stylized design of our model and the indirect nature of the evidence, we do not claim that the actual mechanism is the exact reputational mechanism that our model generates. Yet we have shown that the main alternative, learning by doing, is unlikely to be helpful in this regard. Other candidates such as technological upgrading produce large transitory negative shocks to income growth, but not large positive ones.

Further research should therefore strive to advance and evaluate competing explanations. In this respect, embedding reputational concerns into a standard model of income dynamics appears to be a promising avenue of research. As documented in GKOS, such a richer model would include a mixture of autoregressive processes with different degrees of persistence, along with transitory shocks and abrupt transitions capturing the risk of non-employment. A convincing test for this extended framework would be to assess whether it matches the systematic differences in kurtosis and persistence of earnings growth across income and age groups that cannot be captured by our stylized model. In particular, since the assumption that workers are hand-to-mouth is more likely to hold at the bottom of the earnings distribution, future research should particularly focus on the earnings dynamics of low-income workers.

Finally, the model also offers a new perspective on changes in the volatility of income, namely that a rise in income volatility could be the result of a rise in reputational concerns that, in turn could be called by a reduction in the noise accompany-

ing the relation between effort and output. Agency theory suggests that a reduction in noise raises the incentives to provide effort, and we have shown that this can also help explain the rise in the volatility of earnings.

References

- [1] Arellano, Manuel, Richard Blundell and Stéphane Bonhomme. "Earnings and Consumption Dynamics: A Nonlinear Panel Data Framework." *Econometrica*, vol. 85(3) (2017): 693-734.
- [2] Bar-Isaac, Heski, and Johannes Hörner. "Specialized Careers." *Journal of Economics and Management Strategy*, vol. 23(3) (2014): 601-627.
- [3] Ben-Porath, Yoram. "The Production of Human Capital and the Life Cycle of Earnings." *Journal of Political Economy*, vol. 75(4) (1967): 352-365.
- [4] Board, Simon, and Moritz Meyer-ter-Vehn. "Reputation for Quality." *Econometrica*, vol. 81(6) (2013): 2381-2462.
- [5] Bouvard, Matthieu, and Raphaël Levy. "Horizontal reputation and strategic audience management." Forthcoming in the *Journal of the European Economic Association* (2019).
- [6] Burdett, Kenneth, and Dale T. Mortensen. "Wage differentials, employer size, and unemployment." *International Economic Review* (1998): 257-273.
- [7] Cisternas, Gonzalo. "Two-Sided Learning and the Ratchet Principle." *Review of Economic Studies*, vol. 85(1) (2018): 307-351.
- [8] Dahl, Molly, Thomas DeLeire and Jonathan Schwabish. "Estimates of Year-to-Year Volatility in Earnings and in Household Incomes from Administrative, Survey, and Matched Data." *The Journal of Human Resources*, vol. 46(4) (2011): 750-774.
- [9] Fazzari, Steven, Glenn Hubbard and Bruce Petersen. "Financing Constraints and Corporate Investment." *Brookings Papers on Economic Activity*, vol. 19 (1988): 141-206.

- [10] Green, Edward, and Robert Porter. "Noncooperative collusion under imperfect price information." *Econometrica*, vol. 52(1) (1984): 87-100.
- [11] Greenwood, Jeremy, Zvi Hercowitz, and Gregory W. Huffman. "Investment, capacity utilization, and the real business cycle." *American Economic Review* (1988): 402-417.
- [12] Guvenen, Fatih, Fatih Karahan, Serdar Ozkan and Jae Song. "What Do Data on Millions of U.S. Workers Reveal about Life-Cycle Earnings Dynamics?" NBER Working Papers No. 20913 (2015).
- [13] Holmström, Bengt. "Managerial Incentive Problems: A Dynamic Perspective." *Review of Economic Studies*, vol. 66 (1999): 169-182.
- [14] Hubmer, Joachim. "The Job Ladder and its Implications for Earnings Risk." *Review of Economic Dynamics*, vol. 29 (2018): 172-194.
- [15] Jovanovic, Boyan. "Asymmetric Cycles." *The Review of Economic Studies*, vol. 73(1) (2006): 145-162.
- [16] Jovanovic, Boyan, and Yaw Nyarko. "Learning by doing and the choice of technology." *Econometrica*, vol. 64(6) (1996): 1299-1310.
- [17] Jovanovic, Boyan, and Julien Prat. "Reputation Cycles." NBER Working Paper No. 22703 (2016).
- [18] Kambourov, Gueorgui, and Iouri Manovskii. "Rising Occupational and Industry Mobility in the United States: 1968-1997." *International Economic Review*, vol. 49(1) (2008): 41-79.
- [19] Kaplan, Greg, Giovanni L. Violante, and Justin Weidner. "The Wealthy Hand-to-mouth." *Brookings Papers on Economic Activity*, vol. 45(1) (2014): 77-153.
- [20] Kohlberg, Elon, and Jean-Francois Mertens. "On the Strategic Stability of Equilibria." *Econometrica*, vol. 54(5) (1986): 1003-1037.
- [21] Lange, Fabian. "The Speed of Employer Learning." *Journal of Labor Economics*, vol. 25(1) (2007): 1-35.
- [22] Mailath, George J., and Larry Samuelson. *Repeated Games and Reputations*. Oxford University Press, 2006.

- [23] Parente, Stephen. "Technology Adoption, Learning-by-Doing, and Economic Growth." *Journal of Economic Theory*, vol. 63(2) (1994): 346-369.
- [24] Shimer, Rob. "Convergence in Macroeconomics: The Labor Wedge." *American Economic Journal: Macroeconomics*, vol. 1(1) (2009): 280-97.
- [25] Song, Jae, David Price, Fatih Guvenen, Nicholas Bloom and Till von Wachter. "Firming Up Inequality." NBER Working Papers No. 21199 (2015).
- [26] Ziliak, James, Bradley Hardy and Christopher Bollinger. "Earnings Volatility in America: Evidence from Matched CPS." *Labour Economics*, vol. 18 (2011): 742–754.

A Appendix

A.1 Proof of Proposition 1

We prove that, given the unique consumption function that satisfies the no-profit requirement, the optimality of a conjectured deterministic profile can be established by a first-order condition. In Holmström (1999) sufficiency of the necessary condition immediately followed from the linear impact of actions on payoffs. With concave payoffs, we have to characterize the agent's problem both on and off the equilibrium path. Hence we consider arbitrary strategies and let $\delta_t \equiv a_t - a_t^*$ denote the deviation from conjectured effort at each date. We wish to relate the agent's and market's posteriors about θ , which we denote m^a and m , respectively. Bayes rule implies that

$$m_{t+1}^a - m_{t+1} = \frac{1-\lambda}{\lambda} \sum_{s=0}^t \lambda^{t-s} (x_s^a - x_s) = -\frac{1-\lambda}{\lambda} \sum_{s=0}^t \lambda^{t-s} \delta_s, \quad (26)$$

where $x_t^a \equiv y_t - a_t$. Let Δ_t denote the following weighted mean of past deviations

$$\Delta_t = \lambda \Delta_{t-1} + (1-\lambda) \delta_{t-1} \text{ with } \Delta_0 = 0. \quad (27)$$

It follows from (26) that $m_t^a = m_t - \Delta_t$. We can use this relationship to specify the agent's expectations about the law of motion of beliefs. First notice that $y_t - m_t^a = \theta_t - m_t^a + a_t + \varepsilon_t$. Since $\theta_t - m_t^a$ is normally distributed with mean 0 and variance $\bar{\sigma}_\theta^2$, the innovation process from the agent's standpoint is a normally distributed variable, $u \sim \mathcal{N}(0, \bar{\sigma}_\theta^2 + \sigma_\varepsilon^2)$, such that $y_t = m_t^a + a_t + u_t$. Reinserting this decomposition into the law of motion of market's beliefs, we find that

$$m_{t+1} = \lambda m_t + (1-\lambda) x_t = m_t + (1-\lambda) [\delta_t - \Delta_t + u_t].$$

The normality of u_t implies that the distribution of the posterior $P(m_{t+1})$ satisfies

$$P(m_{t+1} | m_t, \Delta_t, \delta_t) = \Phi \left(\frac{m_{t+1} - [m_t + (1-\lambda)(\delta_t - \Delta_t)]}{(1-\lambda)(\bar{\sigma}_\theta + \sigma_\varepsilon)} \right),$$

where $\Phi(\cdot)$ is the standard normal CDF. Given the per-period utility function $U(m, a^*, \delta) = -\exp(-\gamma[m + a^* - g(a^* + \delta)])$, the agent's value function *on and off* the equilibrium

path is given by the fixed point of the following functional equation

$$\begin{aligned} V_t(m, \Delta) &= \max_{\delta} \left\{ U(m, a_t^*, \delta) + \beta \int V_{t+1}(m', \Delta') dP(m' | m, \Delta, \delta) \right\}, \\ \text{s.t. } m' &= m + (1 - \lambda)[\delta - \Delta + u], \\ \Delta' &= \lambda\Delta + (1 - \lambda)\delta. \end{aligned} \quad (28)$$

Lemma 1 *Assume that the per-period utility function is CARA:*

$$U(c, a) = U(c - g(a)) = -\exp(-\gamma(c - g(a))), \text{ with } \gamma > 0. \quad (29)$$

Given a deterministic profile $\{a_t^\}_{t=0}^\infty$, optimal strategies on and off the equilibrium path are deterministic functions of time and of the stock of past deviations $\Delta_t = \lambda\Delta_{t-1} + (1 - \lambda)(a_t - a_t^*)$, where $\Delta_0 = 0$. In particular, they do not depend on the agent's reputation m_t nor on her output history.*

Proof. The lemma follows from the fact that the value function is of the form

$$V_t(m, \Delta) = \exp(-\gamma m) V_t(0, \Delta). \quad (30)$$

To verify the conjecture, we change the variable of integration of the Bellman equation

$$\begin{aligned} V_t(m, \Delta) &= \max_{\delta} \left\{ -\exp(-\gamma[m + a_t^* - g(a_t^* + \delta)]) \right. \\ &\quad \left. + \beta \int V_{t+1}(m + (1 - \lambda)[\delta - \Delta + u], \Delta') d\Phi\left(\frac{u}{\bar{\sigma}_\theta + \sigma_\epsilon}\right) \right\} \\ &= \exp(-\gamma m) \max_{\delta} \left\{ -\exp(-\gamma[a_t^* - g(a_t^* + \delta)]) \right. \\ &\quad \left. + \beta \int V_{t+1}((1 - \lambda)[\delta - \Delta + u], \Delta') d\Phi\left(\frac{u}{\bar{\sigma}_\theta + \sigma_\epsilon}\right) \right\} \\ &= \exp(-\gamma m) V_t(0, \Delta). \end{aligned}$$

The second equality follows reinserting our guess (30) and using the fact that a_t^* is deterministic and thus independent of m , while the third equality follows directly from the Bellman equation (28) defining V . Since the functional equation is a contraction mapping, the value function is unique and its policy function $\delta(\Delta, t)$ identical across all market's beliefs m . Hence, the agent chooses a sequence of deterministic actions that only depend on the cumulative stock of past deviations Δ and potentially on t . In other words, optimal policies are deterministic *on and off* the equilibrium path. ■

We are now in a position to prove Proposition 1. According to Lemma 1, there is no loss of generality in focusing on sequences of deterministic actions. Let $\mathbf{a} \equiv \{a_0, a_1, \dots\}$ denote a sequence of actions from date 0 onwards. Given that effort in each period belongs to the feasibility set $\mathcal{A} \subseteq \mathbb{R}_+$, sequences are element of the infinite Cartesian product $\mathcal{A}^\infty = \{(a_t)_{t=0}^\infty \mid a_t \in \mathcal{A} \text{ for all } t = 0, 1, \dots\}$. The market belief can be computed using the Bayesian map

$$m_{t+1}(\mathbf{a}, \mathbf{x}^a \mid m_0) = \lambda^t m_0 + \frac{1-\lambda}{\lambda} \sum_{s=0}^t \lambda^{t-s} (x_s^a - a_s^* + a_s), \quad (31)$$

where $\mathbf{x}^a \equiv \{x_t\}_{t=0}^\infty$. We use these notations to define the agent's utility as a functional in the space of action sequences

$$V_0(0, 0) = \max_{\mathbf{a} \in \mathcal{A}^\infty} \mathcal{U}(\mathbf{a}) = \max_{\mathbf{a} \in \mathcal{A}^\infty} E \left[- \sum_{t=0}^{\infty} \beta^t \exp(-\gamma[m_t(\mathbf{a}) + a_t^* - g(a_t)]) \right],$$

where $m_t(\mathbf{a})$ is a shorthand version of $m_t(\mathbf{a}, \mathbf{x}^a \mid m_0)$ introduced in (31).³³ Sufficiency of the necessary condition will follow if we can prove that $\mathcal{U}(a)$ is a concave functional. We establish an even stronger requirement as concavity holds in all periods. To see why, let $p_t(\mathbf{a}) \equiv c_t(\mathbf{a}) - g(a_t)$ denote returns in period t , i.e., consumption net of the disutility of effort. We have

$$\begin{aligned} p_t(\alpha \mathbf{a}^1 + (1-\alpha) \mathbf{a}^2) &= m_t(\alpha \mathbf{a}^1 + (1-\alpha) \mathbf{a}^2) + a_t^* - g(\alpha a_t^1 + (1-\alpha) a_t^2) \\ &\geq \alpha [m_t(\mathbf{a}^1) + a_t^* - g(a_t^1)] + (1-\alpha) [m_t(\mathbf{a}^2) + a_t^* - g(a_t^2)] \\ &= \alpha p_t(\mathbf{a}^1) + (1-\alpha) p_t(\mathbf{a}^2), \end{aligned}$$

for any $\alpha \in (0, 1)$ and all pair of sequences $\mathbf{a}^1, \mathbf{a}^2 \in \mathcal{A}^\infty$. The first equality uses the fact that the market conjecture a_t^* is deterministic. The second inequality holds because the disutility g is convex in a , while the Bayesian map for m_t is linear in $\mathbf{a}$, so that $m_t(\alpha \mathbf{a}^1 + (1-\alpha) \mathbf{a}^2) = \alpha m_t(\mathbf{a}^1) + (1-\alpha) m_t(\mathbf{a}^2)$. Per-period utility is given by $U_t(\mathbf{a}) = (U \circ p_t)(\mathbf{a})$, and since the composition of concave continuous functions is itself concave, the strict concavity of U along with the convexity of $\mathcal{A}$ imply that

$$U_t(\alpha \mathbf{a}^1 + (1-\alpha) \mathbf{a}^2) > \alpha U_t(\mathbf{a}^1) + (1-\alpha) U_t(\mathbf{a}^2),$$

³³Since the innovation process $\mathbf{x}^a$ and the prior m_0 cannot be influenced by the agent, they can be treated as exogenous when studying her optimal control problem.

for any $\alpha \in (0, 1)$ and all pair of sequences $\mathbf{a}^1 \neq \mathbf{a}^2 \in \mathcal{A}^\infty$. We have therefore proven that per-period utility is strictly concave in $\mathbf{a}$, which ensures that

$$\begin{aligned} \mathcal{U}(\alpha \mathbf{a}^1 + (1 - \alpha) \mathbf{a}^2) &= E \left[\sum_{t=0}^{\infty} \beta^t U_t(\alpha \mathbf{a}^1 + (1 - \alpha) \mathbf{a}^2) \right] \\ &> \alpha E \left[\sum_{t=0}^{\infty} \beta^t U_t(\mathbf{a}^1) \right] + (1 - \alpha) E \left[\sum_{t=0}^{\infty} \beta^t U_t(\mathbf{a}^2) \right] \\ &= \alpha \mathcal{U}(\mathbf{a}^1) + (1 - \alpha) \mathcal{U}(\mathbf{a}^2). \end{aligned}$$

Note that the inequality holds because the only source of randomness is the innovation process $\mathbf{x}^a$ defined from the agent's standpoint so that its realizations are by definition independent of $\mathbf{a}$. We can therefore conclude that $\mathcal{U}$ is a *strictly concave* functional in the space of action sequences. Thus any local maximum of $\mathcal{U}$ is also a global maximum. A necessary condition for $\mathcal{U}$ to have a maximum at $\mathbf{a}^*$ is that its Gateaux derivative $d\mathcal{U}(\mathbf{a}^*; h) = 0$ for all $h \in \mathcal{A}^\infty$. The existence of the Gateaux derivative being guaranteed by the convexity of $\mathcal{A}^\infty$, this requirement yields the necessary FOC (11) which, by strict concavity of the objective, is also sufficient.

A.2 Proofs of Propositions 2 to 7

Proof. Proposition 2: We wish to show that the necessary conditions (10) and (11) are equivalent when the agent's utility is CARA. Reinserting (9) into (10), we find that (10) is satisfied when

$$\begin{aligned} g'(a_t^*) &= \frac{1 - \lambda}{\lambda} \sum_{s=t+1}^{\infty} (\beta\lambda)^{s-t} E_t^* \left[\frac{U'(c_s - g(a_s^*))}{U'(c_t - g(a_t^*))} \right] \\ &= \frac{1 - \lambda}{\lambda} \sum_{s=t+1}^{\infty} (\beta\lambda)^{s-t} E_t^* [\exp \gamma (m_t + a_t^* - g(a_t^*) - (m_s + a_s^* - g(a_s^*)))]. \end{aligned} \quad (32)$$

The posterior m cancels out from the expectation because

$$m_{t+1} = \lambda m_t + (1 - \lambda) x_t = m_t + (1 - \lambda) (\theta_t - m_t + \varepsilon_t),$$

so that

$$\begin{aligned} E_t^* [\exp(\gamma(m_t - m_s))] &= E_t^* \left[\exp \left(-\gamma(1 - \lambda) \sum_{i=0}^{s-1} (\theta_{t+i} - m_{t+i} + \varepsilon_{t+i}) \right) \right] \\ &= \exp \left(\frac{\gamma^2(1 - \lambda)^2(\sigma_\varepsilon^2 + \bar{\sigma}_\theta^2)(s - t)}{2} \right). \end{aligned} \quad (33)$$

The first equality holds because m_t is a martingale, while the second equality holds because (i) $\theta_t \sim \mathcal{N}(m_t, \bar{\sigma}_\theta^2)$, and (ii) ε_t is a normally distributed noise with variance σ_ε^2 . The expression above can be simplified using the definition of $\bar{\sigma}_\theta^2$ in (6) and the fact that $\lambda = \bar{\sigma}_\theta^2 / (\sigma_\varepsilon^2 + \bar{\sigma}_\theta^2)$, hence

$$\begin{aligned} (1 - \lambda)^2(\sigma_\varepsilon^2 + \bar{\sigma}_\theta^2) &= \frac{\bar{\sigma}_\theta^4}{\sigma_\varepsilon^2 + \bar{\sigma}_\theta^2} = \frac{4\sigma_\varepsilon^4 \left(\sqrt{1 + \frac{4\sigma_\varepsilon^2}{\sigma_\nu^2}} - 1 \right)^{-2}}{\sigma_\varepsilon^2 \left[1 + 2 \left(\sqrt{1 + \frac{4\sigma_\varepsilon^2}{\sigma_\nu^2}} - 1 \right)^{-1} \right]} \\ &= \frac{4\sigma_\varepsilon^2}{\left(\sqrt{1 + \frac{4\sigma_\varepsilon^2}{\sigma_\nu^2}} - 1 \right) \left(\sqrt{1 + \frac{4\sigma_\varepsilon^2}{\sigma_\nu^2}} + 1 \right)} = \sigma_\nu^2. \end{aligned}$$

Simplifying the expectation in (33) and replacing the resulting expression into (32), we finally obtain

$$\begin{aligned} g'(a_t^*) &= \frac{1 - \lambda}{\lambda} \sum_{s=t+1}^{\infty} (\beta\lambda)^{s-t} E_t^* \left[\frac{U'(c_s - g(a_s^*))}{U'(c_t - g(a_t^*))} \right] \\ &= \frac{1 - \lambda}{\lambda} \sum_{s=t+1}^{\infty} (\beta\lambda)^{s-t} \exp \left(\frac{\gamma^2 \sigma_\nu^2 (s - t)}{2} \right) \exp(\gamma[a_t^* - g(a_t^*) - (a_s^* - g(a_s^*))]), \end{aligned}$$

which is indeed equivalent to (11) since $\beta^{s-t} \exp(\gamma^2 \sigma_\nu^2 (s - t) / 2) = r^{s-t}$. ■

Proof. Propositions 3 and 4: Since Proposition 3 is a subcase of Proposition 4, we directly focus on the general problem where $\sigma_\varepsilon > 0$, so that output is noisy. We are interested in 2-period cycles where $a_t^* = \bar{a}$ when t is even, and $a_t^* = 0$ when t is odd. Using the notation

$$S(a) \equiv \exp(\gamma(1 - \kappa)a), \quad (34)$$

and taking into account the feasibility constraint $a \in [0, \bar{a}]$, we find that the Kuhn-

Tucker conditions resulting from the incentive compatibility constraint (10) are

$$\kappa \leq \frac{1-\lambda}{\lambda} \left(\sum_{i \in \{1,3,5,\dots\}} (r\lambda)^i S(\bar{a}) + \sum_{i \in \{2,4,6,\dots\}} (r\lambda)^i \right) \quad \text{if } t \text{ is even so that } a_t^* = \bar{a}, \quad (35)$$

and

$$\kappa \geq \frac{1-\lambda}{\lambda} \left(\sum_{i \in \{1,3,5,\dots\}} (r\lambda)^i S(\bar{a})^{-1} + \sum_{i \in \{2,4,6,\dots\}} (r\lambda)^i \right) \quad \text{if } t \text{ is odd so that } a_t^* = 0. \quad (36)$$

These two conditions can be re-written as

$$\frac{1}{S(\bar{a})} \leq f(\kappa, r, \lambda) \leq S(\bar{a}), \quad \text{where } f(\kappa, r, \lambda) \equiv \kappa \left[\frac{1 - (r\lambda)^2}{r(1-\lambda)} \right] - r\lambda. \quad (37)$$

Let $r_1 \equiv \kappa / [\kappa + (1-\lambda)(1-\kappa)]$ and $r_0 \equiv \sqrt{\kappa / [\lambda(1-\lambda)(1-\kappa)]} > r_1$, so that $f(\kappa, r_1, \lambda) = 1$ and $f(\kappa, r_0, \lambda) = 0$. Since $f(\kappa, r, \lambda)$ is strictly decreasing in r with $\lim_{r \rightarrow 0} f(\kappa, r, \lambda) = +\infty$, there exists a unique $\underline{r} < r_1$ such that $f(\kappa, \underline{r}, \lambda) = S(\bar{a}) > 1$. The strict monotonicity of $f(\kappa, r, \lambda)$ with respect to r also implies that there exists a unique $\bar{r} \in (r_1, r_0)$ such that $f(\kappa, \bar{r}, \lambda) = 1/S(\bar{a})$. It is easily verified that the two inequalities in (37) are satisfied, and thus cycles are sustainable, whenever $r \in (\underline{r}, \bar{r})$.

Cycles are welfare improving when the rest point $\hat{a} = 0$. To show that this is the case when $r \in (\underline{r}, r_1)$, set $a_s^* = \hat{a}$ for all $s \geq t$ in (11). The resulting incentive constraint reads $\kappa = r(1-\lambda + \lambda\kappa)$, a requirement that cannot be satisfied since $r < r_1$ implies that $\kappa > r(1-\lambda)/(1-r\lambda)$. Hence any $\hat{a} > 0$ cannot be incentive compatible as the agent would like to deviate by investing less than $\hat{a}$, as shown in Fig. 10. However, such deviations are not feasible when $a_t^* = 0$ for all t , and so $\hat{a} = 0$ is indeed the only rest point.

Finally, we consider the limit cases where $\sigma_\varepsilon = 0$. Then λ is also equal to 0 and so $r_1 = \kappa$. Moreover since $f(\kappa, r, 0) = \kappa/r$, we have $\underline{r} = \kappa \exp(-\gamma(1-\kappa)\bar{a})$ and $\bar{r} = \kappa \exp(\gamma(1-\kappa)\bar{a})$, as stated in Proposition 3. ■

Proof. Proposition 5: We first prove that cycles can emerge solely if σ_ε does not exceed a certain threshold. The first step consists in rewriting the necessary condition

Figure 10: PHASE PORTRAITS FOR TWO DIFFERENT VALUES OF κ WHEN $g(a) = \kappa a$.

(35) as follows

$$\kappa \leq \frac{1-\lambda}{\lambda} \left(\sum_{i \in \{1,3,5,\dots\}} (r\lambda)^i S(\bar{a}) + \sum_{i \in \{2,4,6,\dots\}} (r\lambda)^i \right) = \frac{r(1-\lambda)}{1-r\lambda} \left(\frac{S(\bar{a})-1}{1+r\lambda} + 1 \right).$$

The expression on the right-hand side is decreasing in λ and converges to 0 as λ goes to 1. Hence there exists a $\bar{\lambda} \in (0, 1)$ such that the necessary condition (35) is not satisfied whenever $\lambda > \bar{\lambda}$. Given that λ is strictly increasing in σ_ε , and that it goes from 0 to 1 as σ_ε goes from 0 to infinity,³⁴ two-period cycles cannot be sustainable whenever σ_ε is greater than a certain value, which we denote by $\bar{\sigma}_\varepsilon$.

(i) Let's now first consider cases where $\kappa \leq r$. If $\lambda = (r - \kappa)/[r(1 - \kappa)]$, the definition of f in (37) implies that $f(\kappa, r, \lambda) = 1$. Given that $S(\bar{a})$ is greater than 1, condition (37) is necessarily satisfied and so $\bar{\lambda} > (r - \kappa)/[r(1 - \kappa)]$. Moreover, it follows from the continuity of f and the proof of Proposition 4 that there exists a $\underline{\lambda} \leq (r - \kappa)/[r(1 - \kappa)]$ such that the necessary condition is sufficient and satisfied for all $\lambda \in [\underline{\lambda}, \bar{\lambda}]$. Given the relationship between σ_ε and λ , $\underline{\lambda}$ defines a unique $\underline{\sigma}_\varepsilon$ such that two-period cycles are sustainable for all $\sigma_\varepsilon \in [\underline{\sigma}_\varepsilon, \bar{\sigma}_\varepsilon]$.

(ii) We now consider cases where $\kappa > r$. The necessary condition is satisfied when

³⁴Remember that $\lambda = 1 - \frac{\sigma_\varepsilon^{-2}}{\bar{\sigma}_\theta^{-2} + \sigma_\varepsilon^{-2}} = 1 - 2 \left(\sqrt{1 + 4 \frac{\sigma_\varepsilon^2}{\bar{\sigma}_\theta^2}} + 1 \right)^{-1}$.

$\lambda = 0$ if and only if $f(\kappa, r, 0) \leq S(\bar{a})$, i.e., when $\bar{a} \geq \log(\kappa/r)/[\gamma(1-\kappa)]$. Then we also know from the first step of the proof that there exists a $\bar{\lambda}$ such that the necessary condition will be satisfied for all $\lambda \in [0, \bar{\lambda}]$. Moreover, we have shown in Step 3 of the proof of Proposition 4 that the necessary condition is also sufficient whenever the rest point is inefficient, a property that is verified because $\kappa > r$. Finally, using once again the relationship between σ_ε and λ , we conclude that two-period cycles are sustainable if and only if $\sigma_\varepsilon \in [0, \bar{\sigma}_\varepsilon]$. ■

Proof. Proposition 6: We want to characterize deterministic dynamics near the rest point $\hat{a}$. Let $a_t = \varphi(a_{t+1})$ denote the implicit map so that $l(\varphi(a)) = h(a)$, where $l(a) \equiv g'(a)U'(a - g(a))$ and $h(a) \equiv rU'(a - g(a))[1 - \lambda + \lambda g'(a)]$. Differentiating the incentive constraint (13) at the rest point, we find that

$$\varphi'(\hat{a}) \equiv \left. \frac{da_t}{da_{t+1}} \right|_{a_{t+1}=\hat{a}} = \frac{\frac{\lambda g'(\hat{a})}{1-\lambda+\lambda g'(\hat{a})} A_g(\hat{a}) - A_U(\hat{a} - g(\hat{a}))[1 - g'(\hat{a})]}{A_g(\hat{a}) - A_U(\hat{a} - g(\hat{a}))[1 - g'(\hat{a})]}, \quad (38)$$

where $A_U(x) = -U''(x)/U'(x)$ and $A_g(a) = g''(a)/g'(a)$.

The rest point $\hat{a}$ is locally unstable if $|\varphi'(\hat{a})| > 1$. Since the denominator in (38) is always higher than the numerator, the instability condition can be satisfied solely if the numerator is negative. Let us focus first on cases where the denominator is positive while the numerator is negative. Then it is easy to verify that $\varphi'(\hat{a}) < -1$ whenever (19) is satisfied. The other possibility is that both numerator and denominator are negative, so that

$$A_g(\hat{a}) < A_U(\hat{a} - g(\hat{a}))(1 - g'(\hat{a})) \Leftrightarrow \varphi'(\hat{a}) > 1, \quad (39)$$

a condition that is always fulfilled when requirement (19) in Proposition 6 is satisfied. ■

Proposition 7 can be proved in a similar way to Proposition 4. Thus we first use a direct approach by changing variable and defining a new fixed point problem.

Lemma 2 *Let*

$$a(s) \equiv \frac{r(1-\lambda)}{1-(r\lambda)^2} (r\lambda + s),$$

2-period cycles are sustainable when the fixed point problem

$$s = \psi(s) \equiv \frac{U'(c(a(1/s)) - g(a(1/s)))}{U'(c(a(s)) - g(a(s)))} \quad (40)$$

admits a solution $s^ \neq 1$.*

Proof. Lemma 2: We focus on 2-period cycles and denote this period action by a , next period by a' , the period after by a , and so on. In other words, we have the discount factors from today til tomorrow and from tomorrow til the day after, respectively,

$$s \equiv U'(c(a') - g(a'))/U'(c(a) - g(a)), \quad (41)$$

$$s' \equiv 1/s(a, a'). \quad (42)$$

Therefore, if we start at $t = 0$, so that a is the action at $t = 0, 2, 4, 6, \dots$ and a' the action for $t = 1, 3, 5, 7, \dots$, then the incentive constraint (10) is satisfied when (a, a') solve the following two equations

$$a = \frac{1-\lambda}{\lambda} \left(\sum_{t=1,3,5,7,\dots}^{\infty} (r\lambda)^t s + \sum_{t=2,4,6,8,\dots}^{\infty} (r\lambda)^t \right) = \frac{1-\lambda}{\lambda} \left(\frac{r\lambda s}{1-(r\lambda)^2} + \frac{(r\lambda)^2}{1-(r\lambda)^2} \right),$$

$$a' = \frac{1-\lambda}{\lambda} \left(\sum_{t=1,3,5,7,\dots}^{\infty} (r\lambda)^t s' + \sum_{t=2,4,6,8,\dots}^{\infty} (r\lambda)^t \right) = \frac{1-\lambda}{\lambda} \left(\frac{r\lambda s'}{1-(r\lambda)^2} + \frac{(r\lambda)^2}{1-(r\lambda)^2} \right).$$

These simplify to

$$a = \frac{1-\lambda}{1-(r\lambda)^2} r(r\lambda + s), \quad (43)$$

$$a' = \frac{1-\lambda}{1-(r\lambda)^2} r(r\lambda + s'). \quad (44)$$

Thus there are 4 equations (41), (42), (43), and (44), and 4 unknowns, (a, a', s, s') . One solution is $(\hat{a}, \hat{a}, 1, 1)$ where $\hat{a} = (1-\lambda)/(r^{-1}-\lambda)$, which is a version of Holmström's (1999) Proposition 1. Now let's treat s as a parameter to begin with. The action a as a function of s is given by (43) and (44). Thus our problem is equivalent to looking for a fixed point in s of the function $\psi(s)$ defined in (40). ■

Proof. Proposition 7: Since $U(c) = -\exp(-\gamma c)$, it follows from the definition in (40) of $\psi(\cdot)$ that $\psi(1) = 1$. Let $\bar{s} \equiv 2\frac{1-(r\lambda)^2}{r(1-\lambda)} - r\lambda$, since $c(\bar{s}) = 0$ and $c(1/\bar{s}) \in (0, 1)$, we have

$\psi(\bar{s}) < 1$. By continuity of the mapping $\psi(\cdot)$, there will be a fixed point $s \in (1, \bar{s})$ if $\psi'(1) > 1$. Differentiating $\psi(\cdot)$, we obtain

$$\psi'(s) = \gamma \left(1 - a(s) + \frac{1 - a(1/s)}{s^2} \right) \frac{r(1 - \lambda)}{1 - (r\lambda)^2} \exp \left(\gamma \left[a(s) - \frac{a(s)^2}{2} - \left(a(1/s) - \frac{a(1/s)^2}{2} \right) \right] \right),$$

and so

$$\psi'(1) > 1 \Leftrightarrow 2\gamma [1 - a(1)] \frac{r(1 - \lambda)}{1 - (r\lambda)^2} > 1.$$

This expression can be simplified as $a(1) = \hat{a} = (1 - \lambda)/(r^{-1} - \lambda)$ so that

$$\frac{r(1 - \lambda)}{1 - (r\lambda)^2} = \frac{a(1)}{1 + r\lambda}. \quad (45)$$

Reinserting this equality into the previous equation we get

$$\psi'(1) > 1 \Leftrightarrow 2\gamma [1 - a(1)] a(1) > 1 + r\lambda,$$

which implies in turn that

$$\left| \frac{\frac{r\lambda}{a(1)} - \gamma(1 - a(1))}{\frac{1}{a(1)} - \gamma(1 - a(1))} \right| > 1.$$

This allows us to conclude that the rest point is unstable since equation (38) with quadratic costs and CARA utility reads

$$\varphi'(a(1)) = \frac{da_t}{da_{t+1}} \Big|_{a_{t+1}=a(1)} = \frac{\frac{1}{\frac{1-\lambda}{\lambda} + a(1)} - \gamma(1 - a(1))}{\frac{1}{a(1)} - \gamma(1 - a(1))} = \frac{\frac{r\lambda}{a(1)} - \gamma(1 - a(1))}{\frac{1}{a(1)} - \gamma(1 - a(1))},$$

where the last equality follows from expression of the rest-point $a(1) = (1 - \lambda)/(r^{-1} - \lambda)$. ■

A.3 Asymmetric cycles.

Proposition 8 derives conditions under which cycles of period 3 can be constructed using the partial equilibrium model with linear costs described in Subsection 4. It focuses on asymmetric cycles with protracted booms and sudden busts. The logic of the proof can be applied to study reverse cases and to show that cycles with pro-

tracted slumps and sudden booms can be sustained when $\kappa > (1 - \lambda)/(r^{-1} - \lambda)$. Furthermore, it is cumbersome but relatively straightforward to extend the proof so as to construct cycles with more than 3 states.

Proposition 8 *Consider the partial equilibrium model with CARA utility and linear costs, i.e., $g(a) = \kappa a$. Assume that κ satisfies the compatible conditions*

$$\frac{1 - \lambda}{r^{-1} - \lambda} > \kappa > \frac{r(1 - \lambda)}{1 - (r\lambda)^3} \left[\frac{r^2\lambda(1 - \lambda + \lambda\kappa)}{\kappa} + (r\lambda)^2 \right]. \quad (46)$$

Then there exists a unique $\bar{a}_{min} > 0$ such that deterministic cycles of period three, where

$$a_t^* = \begin{cases} \bar{a} & \text{when } t = \{0, 3, 6, \dots\} \\ 0 & \text{when } t = \{1, 4, 7, \dots\} \\ \tilde{a} \in (0, \bar{a}) & \text{when } t = \{2, 5, 8, \dots\}, \end{cases}$$

are sustainable whenever $\bar{a} \geq \bar{a}_{min}$.

Proof. Proposition 8: We propose a constructive proof. We study each action in turn and prove their incentive compatibility

1. $a_t^* = \bar{a}$: Since we are focusing on cases where $\kappa < r(1 - \lambda)/(1 - r\lambda)$, we have

$$\kappa \leq \frac{1 - \lambda}{\lambda} \sum_{i=\{1,2,\dots\}}^{\infty} (r\lambda)^i < \frac{1 - \lambda}{\lambda} \sum_{i=\{1,2,\dots\}} (r\lambda)^i S(\bar{a} - a_{t+i}^*). \quad (47)$$

The last inequality follows from the definition of $S(\cdot)$ in (34) because $\bar{a} \geq a_{t+i}^*$ for all i , and so $S(\bar{a} - a_{t+i}^*) \geq 1$, with strict inequality for some i . Since the last expression in (47) measures the discounted returns from effort at date t , the feasibility constraint binds and $a_t^* = \bar{a}$ is indeed incentive compatible.

2. $a_t^* = 0$: At the lower-bound of the feasibility set, costs must exceed returns so

that

$$\begin{aligned}
\kappa &\geq \frac{1-\lambda}{\lambda} \sum_{i=1,2,\dots}^{\infty} (r\lambda)^i S(-a_{t+i}^*) \\
&= S(-\tilde{a}) r \left(1 - \lambda + \lambda \left[\frac{1-\lambda}{\lambda} \sum_{i=\{2,3,\dots\}} (r\lambda)^{i-1} S(\tilde{a} - a_{t+i}^*) \right] \right) \\
&= S(-\tilde{a}) r (1 - \lambda + \lambda\kappa).
\end{aligned}$$

The last equality follows from the fact that the incentive constraint must hold with equality in the next period because $a_{t+1}^* = \tilde{a} \in (0, \bar{a})$. The condition is satisfied whenever

$$\tilde{a} \geq \tilde{a}_{min} \equiv \frac{-\log\left(\frac{\kappa}{r(1-\lambda+\lambda\kappa)}\right)}{\gamma(1-\kappa)} > 0. \quad (48)$$

3. $a_t^* = \tilde{a} \in (0, \bar{a})$: First, we assume that condition (48) holds as an equality and we show that there exists a unique value of $\bar{a}$ which renders $\tilde{a}_{min}$ incentive compatible. Since $\tilde{a}$ is interior, the incentive constraint has to hold exactly, i.e.,

$$\begin{aligned}
\kappa &= \frac{1-\lambda}{\lambda} \sum_{i=1,2,\dots}^{\infty} (r\lambda)^i S(\tilde{a} - a_{t+i}^*) \\
&= S(\tilde{a} - \bar{a}) r \left(1 - \lambda + \lambda \left[\frac{1-\lambda}{\lambda} \sum_{i=\{2,3,\dots\}} (r\lambda)^{i-1} S(\bar{a} - a_{t+i}^*) \right] \right) \\
&> S(\tilde{a} - \bar{a}) r (1 - \lambda + \lambda\kappa) > S(\tilde{a} - \bar{a}) \kappa.
\end{aligned} \quad (49)$$

The first inequality follows from step 1 above, while the second one holds because $\kappa < r(1-\lambda)/(1-r\lambda)$. It shows that $S(\tilde{a} - \bar{a})$ has to be below one, thus requiring that $\bar{a}$ be strictly larger than $\tilde{a}$. Keeping $\tilde{a}$ constant and differentiating (49) with respect to $\bar{a}$, we find that returns are strictly decreasing in $\bar{a}$. Furthermore, if

$$\kappa > \frac{1-\lambda}{\lambda} \left[\frac{(r\lambda)^2 r (1-\lambda+\lambda\kappa)/\kappa + (r\lambda)^3}{1-(r\lambda)^3} \right], \quad (50)$$

there exists a unique value, which we denote $\bar{a}_{min}$, such that $\tilde{a} = \tilde{a}_{min}$ and (49) is

satisfied. To verify that (50) is consistent with $r(1 - \lambda)/(1 - r\lambda) > \kappa$, notice that

$$\frac{r(1 - \lambda)}{1 - r\lambda} > \frac{1 - \lambda}{\lambda} \left[\frac{(r\lambda)^2 r (1 - \lambda + \lambda\kappa) / \kappa + (r\lambda)^3}{1 - (r\lambda)^3} \right] \Leftrightarrow 1 > r\lambda \left[\frac{r(1 - \lambda + \lambda\kappa)}{\kappa} - 1 \right].$$

This inequality is equivalent to $\kappa > r^2\lambda(1 - \lambda) / [1 + r\lambda - (r\lambda)^2]$, which yields a lower-bound that is less than the term on the right-hand side of (46) as well as to $r(1 - \lambda)/(1 - r\lambda)$. Hence the two conditions in (46) are compatible since they define a non-empty interval. As with period-2 cycles, low levels of effort can never be incentive compatible when costs are too small.

We still have to prove that our cycles are sustainable when $\bar{a} > \bar{a}_{min}$. Differentiating (49) with respect to both $\bar{a}$ and $\tilde{a}$, one finds that $\partial\tilde{a}/\partial\bar{a} \in (0, 1)$. The derivative being positive, $\tilde{a} > \tilde{a}_{min}$ if $\bar{a} > \bar{a}_{min}$, and the condition (48) for incentive compatibility of $a_t^* = 0$ is satisfied. Furthermore, the derivative being smaller than one ensures that $\tilde{a}$ remains within the interior of the feasible set as $\bar{a}$ increases. ■

A.4 Stochastic cycles.

Proposition 9 : *Let the effort pair $(\underline{a}(s), \bar{a}(s))$ be defined as follows*

$$\underline{a}(s) = r(1 - \lambda)\alpha \left[\frac{r\lambda\alpha + s + (1 - \alpha)(1 - r\lambda(1 - \alpha))/\alpha}{(1 - r\lambda(1 - \alpha))^2 - (r\lambda\alpha)^2} \right], \quad (51)$$

$$\bar{a}(s) = r(1 - \lambda)\alpha \left[\frac{r\lambda\alpha + s^{-1} + (1 - \alpha)(1 - r\lambda(1 - \alpha))/\alpha}{(1 - r\lambda(1 - \alpha))^2 - (r\lambda\alpha)^2} \right]. \quad (52)$$

If the fixed point problem

$$s = \frac{U'(c(\bar{a}(s)) - g(\bar{a}(s)))}{U'(c(\underline{a}(s)) - g(\underline{a}(s)))} \quad (53)$$

admits a solution $s^ \neq 1$, stochastic cycles where effort switches with probability $\alpha \in (0, 1)$ from $\underline{a}(s)$ to $\bar{a}(s)$, or from $\bar{a}(s)$ to $\underline{a}(s)$, are sustainable*

Proof. Proposition 9: Let $s = U'(c(\bar{a}) - g(\bar{a}))/U'(c(\underline{a}) - g(\underline{a}))$ denote the discount factor when effort switches from $\underline{a}$ to $\bar{a}$. Then the recursive incentive constraint (13) is satisfied when the effort pair $(\underline{a}, \bar{a})$ solves the following system of equations

$$\begin{aligned} \underline{a} &= r[(1 - \alpha)(1 - \lambda + \lambda\underline{a}) + \alpha s(1 - \lambda + \lambda\bar{a})], \\ \bar{a} &= r[(1 - \alpha)(1 - \lambda + \lambda\bar{a}) + \alpha s^{-1}(1 - \lambda + \lambda\underline{a})]. \end{aligned}$$

Combining the two equations above yields (51) and (52). As expected, when $\alpha = 1$, we recover the expressions (43) and (44) in Lemma 2. ■