Spence meets Holmström: Luck and repetition in signalling

Sander Heinsalu *

September 23, 2015

Abstract

This paper studies repeated costly signalling when luck matters for the outcome. Benefit is obtained from the belief of the market, not directly from the effort or the signal. Nonstationary environments are allowed.

In the unique equilibrium in which effort is affine in type, the more the current cost of effort varies in type, the smaller the effort of the lowest types and the greater that of the highest. The effort and payoff of intermediate types generically moves in the same direction as their cost and marginal cost. Greater dependence of future costs on type reduces effort for all types.

The framework extends to productive effort, human capital accumulation and exogenous information revelation.

Keywords: Signalling, dynamic games, stochastic games, incomplete information.

JEL classification: C73, D82, D83.

1 Introduction

Acquiring education to signal one's intelligence takes time. The studying effort translates into education outcomes only stochastically (grades are partly due to luck). Intelligence (IQ) is normally distributed (the famous 'bell curve' of Herrnstein and Murray (1994)). The grade distribution is also often normal.[1] To describe such situations, a model of repeated noisy signalling with normally

*School of Economics, University of Queensland.
Email: s.heinsalu@uq.edu.au, website: http://sanderheinsalu.com/, address: Level 6, Colin Clark Building (39), St Lucia, Brisbane, QLD 4072, Australia.
The author thanks seminar audiences at the University of Tartu, the Australasian Economic Theory Workshop, the Australian National University, the University of Sydney, the 30th Annual Congress of the European Economic Association and the Collegio Carlo Alberto for their comments. All remaining errors are the author's.

[1] The grade is often the sum of marks for many questions. If these marks are independent conditional on knowledge, then the distribution of their sum converges to normal as the number of questions increases.

distributed types and signals is proposed. Repetition means that signalling effort is chosen on multiple occasions, not just once. Noise means the effort is imperfectly observable.

Other examples of repeated noisy signalling are a politician practising speeches in order to appear competent in the run-up to an election, a lobbyist exerting effort to shift public opinion to induce politicians to decide favourably, or a manager artificially inflating financial results every reporting period to get a bonus based on the share price.

The repeated noisy signalling model proposed in this paper is technically simple, tractable even when the environment is nonstationary, and yields closed form comparative statics. Some of the results are as expected, e.g. signalling effort increases in the precision of the signal and the variance of the type distribution. Some are more surprising—for example, as the dependence of cost on type increases, some of the types whose cost rises raise their effort. An increased influence of type on cost has qualitatively different effects on equilibrium effort depending on whether the influence increases at present or in the future.

Time is discrete and the horizon finite or infinite. There is a sender (employee) and a competitive market (employers). The sender has an IQ level (type) drawn from a normal distribution. Type determines the cost of effort. The wage the market would be willing to offer the sender under perfect information is affine in type and effort. The sender privately observes the type and chooses his effort level every period. The market has a common prior belief about the sender's IQ and observes in every period a normally distributed signal (test result) with mean equal to the effort. In some (ex ante known) periods, the sender receives a benefit, which is determined by the posterior belief of the market. Attention is restricted to *affine* equilibria (the sender's effort is affine in type).

In the unique affine equilibrium, the greater the current effect of IQ on the cost of effort, the less effort the least intelligent and the more the most intelligent exert. Intuitively, the cost of the lowest types rises, so they exert less effort. Similarly, the highest types exert more effort because their cost falls. However, generically there are types whose effort moves in the same direction as their cost and marginal cost. The reason is that a greater dependence of cost on type tends to make the benefit respond more to the signal. This increases the payoff to signalling, which may outweigh the increased cost.

The more the future cost of effort is influenced by IQ, the lower the current effort of all types and the higher the future effort. Signalling efforts in different periods are substitutes. This is due to the increased precision of the future belief of the employers when the future dependence of cost on IQ increases. A more precise belief responds less to effort, reducing the present benefit of signalling.

The effort of all types increases in the precision of the signal. The intuition is that less noise implies a greater benefit of signalling. Effort increases in the variance of the type distribution, because the reward the market is willing to offer to different types becomes more dispersed. There is thus a greater incentive to masquerade as a higher type, which must be counteracted by greater, costlier effort of the higher types. Effort is independent of the mean of the type distribution, because the incentive to signal only comes from the wage difference between types, not the level of wages.

The next section introduces the model and Section 3 discusses equilibria. Comparative statics are calculated in Section 4. Section 5 extends the model to exogenously and endogenously changing type, exogenous information revelation, productive effort (similar to the career concerns model of Holmström (1999)) and multiple senders. Section 6 discusses the related literature and Section 7 concludes.

2 Model

The players are a sender and a competitive market. Time is discrete, indexed by t. The horizon length is $T \in \mathbb{N} \cup \{\infty\}$. The sender's type is $\theta \in \mathbb{R}$. The type distribution is normal: $p_0 = N(\mu_\theta, \tau_\theta)$, where $\tau_X = \frac{1}{\sigma_X^2}$ denotes the precision and σ_X^2 the variance of random variable X.

Each period, the sender chooses effort $e \in \mathbb{R}$.[2] The action and utility of the market are not modelled explicitly. Instead, the sender is assumed to derive a benefit directly from the belief $p^T = (p_1, \ldots, p_T)$ of the market, where $p_t \in \Delta(\mathbb{R})$ is the belief at the end of period t (when t signals have been observed). The assumption that benefit depends directly on belief follows Spence (1973). It can be microfounded similarly to Cho and Kreps (1987).

The market does not observe the sender's effort e, past or present. Only a noisy signal $s_t = e_t + \epsilon_t$ of the sender's effort is publicly observed, where $\epsilon_t \sim N(0, \tau_{\epsilon t})$ is independent over time. Based on the signal, the market will update its belief about the type of the sender. Denote the mean and precision of the posterior belief p_t by $\mu_{\theta t}$ and $\tau_{\theta t}$ respectively.

The set of signal histories of length t is $\mathbb{R}^t$. A generic history of length t is denoted $s^t = (s_1, \ldots, s_t)$. Due to the noisy observation of effort, all signal histories are on the equilibrium path and refinements are not applicable. A public strategy of the sender is

$$e^T = (e_1(\theta), \ldots, e_T(\theta))_{\theta \in \mathbb{R}} : \mathbb{R} \times \cup_{t=0}^{T-1} \mathbb{R}^t \to \mathbb{R}.$$

The effort path of a single type θ is denoted $e^T(\theta) = (e_1(\theta), \ldots, e_T(\theta))$.

Public strategies are w.l.o.g., because payoffs from the current period on depend only on the public signals generated and efforts chosen. The sender gains nothing by conditioning the current action on past ones.

For any strategy expected from the sender, the sender's best response is pure, because convex cost implies that for any mixed effort, the same signal distribution can be generated at lower cost by a pure effort level.

The sender's stage game payoff at time t is

$$u_{St}(\theta, e_t, p_t) = b_t \mu_{\theta t} - c_t (e_t - \alpha_t \theta),$$

[2] Negative efforts are possible, but their probability in the equilibrium studied later can be made arbitrarily small by making μ_θ positive and large enough. Alternatively, if one wishes effort to be positive, e can be thought of as the log of effort.

with $b_t \geq 0$, $c_t \in C^2$, $c_t \geq 0$, $c_t' > 0$, $\lim_{z\to\infty} c_t'(z) = \infty$, $c_t'' > 0$, $\alpha_t \geq \underline{\alpha} > 0$. Assuming a cost function of the form $c(e - \alpha\theta)$ ensures the existence of an equilibrium where effort is affine in type. This assumption on cost is with loss of generality, but makes the model tractable. The sender is risk-neutral.

Discounting can be incorporated in b_t and $c_t(\cdot)$, so a separate discount factor is not specified. The sender's overall payoff is

$$U_S(\theta, e^T(\theta), p^T) = \sum_{t=1}^{T} \left[b_t \mu_{\theta t} - c_t\left(e_t(\theta) - \alpha_t\theta\right)\right].$$

To ensure the expectation of $U_S(\theta, e^T(\theta), p^T)$ is finite when $T = \infty$, assume $\sum_{t=1}^{T} b_t < \infty$. To make the game nontrivial, assume $\sum_{t=1}^{T} b_t > 0$.

By assumption, effort is costly in all periods. Benefit is obtained only in periods in which $b_t > 0$. If the benefit is zero from a certain period on, then no type will exert any effort after that period, so the game effectively ends there. Assume w.l.o.g. that there is a benefit to signalling in the last period ($b_T > 0$) if the horizon is finite ($T < \infty$), and $b_t > 0$ infinitely often if the horizon is infinite.

Definition 1. A public equilibrium consists of a public strategy e^{T*} and a belief process p^T such that

(a) $e^{T*}(\theta) \in \arg\max_{e^T(\theta)} \mathbb{E}_{p_0}^{e^T(\theta)}\left[U_S(\theta, e^T(\theta), p^T)\right]$ given p_0, θ and p^T,

(b) p_t is derived from Bayes' rule given e^{T*} and $s^t = (s_1, \ldots, s_t)$:

$$p_t(\theta) = \frac{\Pr(s^t|e^{T*}(\theta))p_0(\theta)}{\int_{-\infty}^{\infty} \Pr(s^t|e^{T*}(z))p_0(z)dz} \quad \forall\theta,$$

where $\Pr(s^t|e^{T*}(\theta)) = \prod_{n=1}^{t} \frac{\sqrt{\tau_{\epsilon n}}}{\sqrt{2\pi}} \exp\left(-\frac{1}{2}(s_n - e_n^*(\theta))^2\tau_{\epsilon n}\right)$.

The expectation in the definition is over the paths of the belief process given the chosen strategy. Henceforth, public equilibria are simply called equilibria.

3 Equilibrium effort and payoff

If all types are expected to exert the same effort in a given period, then in that period, all types will try to choose the lowest-cost effort. This does not exist if $e \in \mathbb{R}$. If the action set is extended to $\mathbb{R} \cup \{-\infty\}$, then pooling on $e = -\infty$ in all periods is always an equilibrium. The equilibria that feature pooling on $e = -\infty$ after some history disappear if the signal depends on type (for some $\eta > 0$, have $s = e + \epsilon + \eta\theta$). Pooling on any $e > -\infty$ is not possible in any equilibrium after any history, because all types will deviate to arbitrarily low effort when pooling is expected.

Henceforth, the focus is on action set $\mathbb{R}$ and equilibria where effort is affine in type after any history. Only the strategy the market expects is restricted to be affine. The sender is allowed to choose any public strategy in response. It will be shown the unique best response is affine. The existence of equilibria with effort nonlinear in type is an intractable problem in general, even in the one-shot game. It is technically similar to the uniqueness question in the insider trading model of Kyle (1985). With exponential cost, the affine equilibrium can be proved unique in a larger set of strategies.

3.1 Finite horizon

The finite horizon case is solved by backward induction. The final period (one-shot game) is solved first. To simplify notation, the time indices are dropped where no confusion arises.

3.1.1 One-shot signalling

Assume that the strategy the market expects from the sender is such that the final period marginal benefit m does not depend on effort or type. The correctness of this expectation will be verified later. Setting marginal benefit equal to marginal cost yields the best response $e(\theta) = (c')^{-1}(m) + \alpha\theta$. The SOC holds due to convex cost. Corner solutions cannot occur due to $e \in \mathbb{R}$. Therefore in any equilibrium with constant marginal benefit, effort is affine in type and the slope in type is α.

To check that the marginal benefit is indeed constant, the belief updating formula must first be found. From the market's point of view, effort e is a random variable (a function of the random variable θ). Updating uses $s|e \sim N(e, \tau_\epsilon)$ and the prior distribution of e from the market's point of view

$$e \sim N\left((c')^{-1}(m) + \alpha\mu_\theta, \quad \frac{\tau_\theta}{\alpha^2}\right).$$

The updated distribution of e given signal s is

$$\begin{aligned} e|s &\sim N\left(\frac{\tau_e\mu_e + \tau_\epsilon s}{\tau_e + \tau_\epsilon}, \quad \tau_e + \tau_\epsilon\right) \\ &= N\left(\frac{\frac{\tau_\theta}{\alpha^2}\left((c')^{-1}(m) + \alpha\mu_\theta\right) + \tau_\epsilon s}{\frac{\tau_\theta}{\alpha^2} + \tau_\epsilon}, \quad \frac{\tau_\theta}{\alpha^2} + \tau_\epsilon\right). \end{aligned}$$

Since $\theta = \frac{1}{\alpha}e - \frac{(c')^{-1}(m)}{\alpha}$, the updated distribution of θ is

$$\theta|s \sim N\left(\frac{\tau_\theta\left((c')^{-1}(m) + \alpha\mu_\theta\right) + \alpha^2\tau_\epsilon s}{\alpha(\tau_\theta + \alpha^2\tau_\epsilon)} - \frac{(c')^{-1}(m)}{\alpha}, \quad \tau_\theta + \alpha^2\tau_\epsilon\right).$$

The average signal equals the effort, so the expected benefit of effort is

$$b\frac{\tau_\theta\mu_\theta + \alpha\tau_\epsilon e - (c')^{-1}(m)\alpha\tau_\epsilon}{\tau_\theta + \alpha^2\tau_\epsilon}. \tag{1}$$

The marginal benefit of effort is $\frac{b\alpha\tau_\epsilon}{\tau_\theta+\alpha^2\tau_\epsilon} = m$ at every effort and signal for every type. The conjecture that the marginal benefit is independent of effort and type is verified. This marginal benefit corresponds to an equilibrium in which effort in terms of primitives is

$$e^*(\theta) = (c')^{-1}\left(\frac{b\alpha\tau_\epsilon}{\tau_\theta+\alpha^2\tau_\epsilon}\right) + \alpha\theta. \tag{2}$$

The expected payoff under the optimal strategy is

$$b\frac{\tau_\theta\mu_\theta+\alpha^2\tau_\epsilon\theta}{\tau_\theta+\alpha^2\tau_\epsilon} - c\left((c')^{-1}\left(\frac{b\alpha\tau_\epsilon}{\tau_\theta+\alpha^2\tau_\epsilon}\right)\right).$$

It was shown that if the marginal benefit is constant, then effort is affine in type with slope α and if effort is affine in type (whatever the slope), then the marginal benefit is constant. This proves the following.

Lemma 1. *If a strategy of the form $e^*(\theta) = k_1 + k_2\theta$ with $k_2 > 0$ is expected, then any best response has the form $e(\theta) = k_3 + \alpha\theta$.*

It was shown above that an expected strategy of the form $e^*(\theta) = k_3+\alpha\theta$ has a unique best response (2). Uniqueness of affine equilibrium is an implication.

The proof of uniqueness in a larger class of strategies closely follows McLennan et al. (2014) Theorem 1.2. Complex-variable functions are used, so more definitions are needed. A *real entire* function is smooth and coincides on $\mathbb{R}$ with its Taylor series centered at zero. A *region* is an open connected set $D \subseteq \mathbb{C}$. A function is *analytic* on D if it is complex-differentiable at every point in D. An *entire* function is analytic on $\mathbb{C}$. An analytic function is *single-valued* if it has an unambiguously defined maximal analytic continuation. If a real-valued function coincides on $\mathbb{C}$ with its Taylor series centered at zero and is smooth, then it is single-valued.

Proposition 2. *If $c(\cdot) = \exp(\cdot)$ or $(c')^{-1}$ is entire, then in the one-shot signalling game there is only one equilibrium that on some $(x_1, x_2) \subset \mathbb{R}$ coincides with a function that is single-valued on (x_1, x_2).*

Proof. For any strategy e^* that the market expects, the posterior mean $\mathbb{E}[\mu_\theta(s)|e] = \int_{-\infty}^{\infty} \mu_{\theta 1}(s)\frac{1}{\pi\sqrt{2}}\exp\left(-(s-e)^2\tau_\epsilon\right) ds$ that the worker expects is entire as a function of effort. This is proved in McLennan et al. (2014) Theorems 2.1 and 2.2. Note that the posterior mean $\mu_{\theta 1}(s)$ after signal s depends on e^*.

Rearrange the FOC $\frac{\partial\mathbb{E}[\mu_{\theta 1}(s)|e]}{\partial e}|_{e=e^*(\theta)} - c'(e^*(\theta)-\alpha\theta) = 0$ as

$$\frac{1}{\alpha}\left[e^*(\theta) - (c')^{-1}\left(\frac{\partial\mathbb{E}[\mu_{\theta 1}(s)|e]}{\partial e}|_{e=e^*(\theta)}\right)\right] = \theta. \tag{3}$$

The left-hand side (LHS) is entire in $e^*(\theta)$ if $(c')^{-1}\left(\frac{\partial\mathbb{E}[\mu_{\theta 1}(s)|e]}{\partial e}|_{e=e^*(\theta)}\right)$ is. The derivative of an entire function is entire. The composition of entire functions is entire, so if $(c')^{-1}$ is entire, then the LHS is as well. The logarithm of a nowhere zero entire function is entire, so another sufficient condition is $\frac{\partial \mathbb{E}[\mu_{\theta 1}(s)|e]}{\partial e}|_{e=e^*(\theta)} \neq 0$.

To prove $\frac{\partial \mathbb{E}[\mu_{\theta 1}(s)|e]}{\partial e}|_{e=e^*(\theta)} > 0 \; \forall \theta$ in any equilibrium, use the incentive constraints (ICs) that an equilibrium e^* must satisfy: $\mathbb{E}[\mu_{\theta 1}(s)|e_1] - c(e_1 - \theta_1) \geq \mathbb{E}[\mu_{\theta 1}(s)|e_2] - c(e_2 - \theta_1)$ and $\mathbb{E}[\mu_{\theta 1}(s)|e_2] - c(e_2 - \theta_2) \geq \mathbb{E}[\mu_{\theta 1}(s)|e_1] - c(e_1 - \theta_2)$. If $\mathbb{E}[\mu_{\theta 1}(s)|e_1] = \mathbb{E}[\mu_{\theta 1}(s)|e_2]$, then $e_1 = e_2$. Adding the ICs,

$$c(e_1 - \theta_1) + c(e_2 - \theta_2) \leq c(e_2 - \theta_1) + c(e_1 - \theta_2).$$

If $\theta_1 < \theta_2$ and $e_1 \geq e_2$, then $e_1 - \theta_1 > \{e_1 - \theta_2, e_2 - \theta_1\} > e_2 - \theta_2$, so $e_1 - \theta_1, e_2 - \theta_2$ is a mean-preserving spread of $e_1 - \theta_2, e_2 - \theta_1$. The mean is $\frac{1}{2}[e_1 + e_2 - \theta_1 - \theta_2]$. Applying a convex $c(\cdot)$ to a mean-preserving spread, one gets $c(e_1 - \theta_1) + c(e_2 - \theta_2) > c(e_2 - \theta_1) + c(e_1 - \theta_2)$, a contradiction. So $\theta_1 < \theta_2 \Rightarrow e_1 < e_2$. The strict monotone likelihood ratio property (MLRP) of normal signals now implies $e_1 < e_2 \Rightarrow \mathbb{E}[\mu_{\theta 1}(s)|e_1] < \mathbb{E}[\mu_{\theta 1}(s)|e_2]$. Smoothness of $\mathbb{E}[\mu_{\theta 1}(s)|e]$ means that $\frac{\partial \mathbb{E}[\mu_{\theta 1}(s)|e]}{\partial e}|_{e=e^*(\theta)} > 0$.

Proposition 3.2 of McLennan et al. (2014) applies unchanged to 3, proving that if e^* coincides on some (x_1, x_2) with a function that is single-valued on (x_1, x_2), then e^* is affine. There is only one affine equilibrium. □

Having studied the one-shot case, the next step is to examine a multiperiod model. It turns out the equilibrium effort and payoff are similar to the one-shot case.

3.1.2 Arbitrary period

Belief remains normal under Bayesian updating when the effort expected by the market is affine in type. In a multiperiod model with an affine expected strategy $e_t^*(\theta) = k_t + \alpha_t \theta \; \forall t$, the precision of the belief is updated deterministically and independently of type, realized signal or chosen effort (but depending on the expected effort) by the formula $\tau_{\theta,t} = \tau_{\theta,t-1} + \alpha_t^2 \tau_{\epsilon t}$. In terms of the precision of the prior τ_θ, the precision at the end of period t is $\tau_{\theta,t} = \tau_\theta + \sum_{n=1}^{t} \tau_{\epsilon n} \alpha_n^2$.

The marginal benefit of signalling in period t consists of the marginal flow benefit m_t (which comes from shifting the belief at the end of the current period) and the marginal benefit of shifting the starting belief in all future periods. The latter consists of the effect of the mean of the belief on the benefit in each period, multiplied by the effect of the mean at the end of the current period on the appropriate future period's mean, summed and multiplied by the expected change in the mean at the end of the current period induced by higher effort:

$$\frac{\alpha_t \tau_{\epsilon t}}{\tau_{\theta,t}} \left[\frac{b_{t+1} \tau_{\theta,t}}{\tau_{\theta,t+1}} + \sum_{k=2}^{T-t} \frac{b_{t+k} \tau_{\theta,t+k-1}}{\tau_{\theta,t+k}} \prod_{j=1}^{k-1} \frac{\tau_{\theta,t+j-1}}{\tau_{\theta,t+j}} \right] = \alpha_t \tau_{\epsilon t} \sum_{k=1}^{T-t} \frac{b_{t+k}}{\tau_{\theta,t+k}}.$$

The notational convention used is that if $k < y$, then $\prod_{j=y}^{k} x_j = 1$ for any x_j.

Conjecture that the marginal flow benefit is constant in effort and type. Define $\gamma_{T-t} = (c_t')^{-1}\left(m_t + \alpha_t\tau_{\epsilon t}\sum_{k=1}^{T-t}\frac{b_{t+k}}{\tau_{\theta,t+k}}\right)$. Equating marginal benefit and marginal cost, the sender's best response is $e_t(\theta) = \gamma_{T-t} + \alpha_t\theta$.

The updated distribution of effort from the market's viewpoint conditional on signal s_t is

$$e_t|s_t \sim N\left(\frac{\tau_{\theta,t-1}\left[\gamma_{T-t} + \alpha_t\mu_{\theta,t-1}\right] + \alpha_t^2\tau_{\epsilon t}s_t}{\tau_{\theta,t-1} + \alpha_t^2\tau_{\epsilon t}}, \quad \frac{\tau_{\theta,t-1}}{\alpha_t^2} + \tau_{\epsilon t}\right).$$

Using $\theta = \frac{1}{\alpha_t}e_t - \frac{\gamma_{T-t}}{\alpha_t}$, the updated distribution of θ is

$$\theta|s_t \sim N\left(\frac{\tau_{\theta,t-1}\mu_{\theta,t-1} + \alpha_t\tau_{\epsilon t}(s_t - \gamma_{T-t})}{\tau_{\theta,t-1} + \alpha_t^2\tau_{\epsilon t}}, \quad \tau_{\theta,t-1} + \alpha_t^2\tau_{\epsilon t}\right), \tag{4}$$

so the derivative of the mean of the belief after period t w.r.t. the mean before t is $\frac{\tau_{\theta,t-1}}{\tau_{\theta,t-1}+\alpha_t^2\tau_{\epsilon t}}$. This is used in the derivation of the marginal benefit of shifting the mean of future belief.

The expected flow benefit in period t from effort $e_t(\theta)$ is

$$b_t\left[\frac{\tau_{\theta,t-1}\left(\gamma_{T-t} + \alpha_t\mu_{\theta,t-1}\right) + \alpha_t^2\tau_{\epsilon t}e_t(\theta)}{\alpha_t(\tau_{\theta,t-1} + \alpha_t^2\tau_{\epsilon t})} - \frac{\gamma_{T-t}}{\alpha_t}\right].$$

The marginal flow benefit is $m_t = \frac{b_t\alpha_t\tau_{\epsilon t}}{\tau_{\theta,t-1}+\alpha_t^2\tau_{\epsilon t}}$. The total marginal benefit is the sum of the marginal flow benefit and the payoff from shifting the mean of the belief in the future

$$\frac{b_t\alpha_t\tau_{\epsilon t}}{\tau_{\theta,t-1} + \alpha_t^2\tau_{\epsilon t}} + \alpha_t\tau_{\epsilon t}\sum_{k=1}^{T-t}\frac{b_{t+k}}{\tau_{\theta,t+k}} = \alpha_t\tau_{\epsilon t}\sum_{k=0}^{T-t}\frac{b_{t+k}}{\tau_{\theta,t+k}},$$

which is constant in effort and type, as conjectured. The optimal effort is

$$e_t^*(\theta) = (c_t')^{-1}\left(\alpha_t\tau_{\epsilon t}\sum_{n=t}^{T}\frac{b_n}{\tau_{\theta,t-1} + \sum_{k=t}^{n}\tau_{\epsilon k}\alpha_k^2}\right) + \alpha_t\theta. \tag{5}$$

In any finite horizon signalling game there is by backward induction a unique affine informative equilibrium.

As (5) shows, optimal effort does not depend on the mean of the belief. Since the precision of the belief changes deterministically over time, effort does not depend on the sender observing the signal. Therefore the model with private monitoring where only the market sees the signal (for example, employers see reference letters that the employee does not) has the same unique affine informative equilibrium. The same holds for any intermediate information structure where the sender sees some noisy signal of the signal (e.g. the interviewer's questions convey information about the content of the reference letters).

3.2 Infinite horizon

Taking the limit as $T \to \infty$ of finite-horizon games and their respective equilibria, the effort in the infinite-horizon game is

$$e_t^*(\theta) = (c_t')^{-1}\left(\alpha_t \tau_{\epsilon t} \sum_{k=0}^{\infty} \frac{b_{t+k}}{\tau_{\theta,t+k}}\right) + \alpha_t \theta, \tag{6}$$

with $\tau_{\theta,t+k} = \tau_{\theta,t} + \sum_{n=t+1}^{k} \tau_{\epsilon n}\alpha_n^2$. That this is an equilibrium can be checked by following the same reasoning as in the finite horizon case—assume the marginal benefit is independent of effort or type, take the FOC, update the belief and verify the assumption on marginal benefit. Effort is finite for every type after every history, because $\sum_{t=1}^{\infty} b_t < \infty$ by assumption. If there is some $\delta \in (0,1)$ s.t. $c_t = \delta^t c$, $b_t = \delta^t b$ and $\alpha_t = \alpha$ for all t, then the discounted sum of efforts is also finite for every type and sequence of signals.

There are no other affine equilibria in the infinite horizon game. Together with the preceding paragraph, we have

Proposition 3. *There is a unique affine equilibrium in the infinite horizon game, with strategy given by (6).*

The proof of uniqueness is the same as in the one-shot case: if effort is expected to be affine in type, then the marginal benefit is constant, which implies that the slope of effort in type is α_t in period t. The slope determines the intercept and thus the marginal benefit.

4 Comparative statics

The formulas in this section do not depend on whether the horizon T is finite or infinite.

Effort at t clearly increases in type and the current and future benefit multipliers $b_{t+k},\ k \geq 0$, based on (5). It also increases in the precision of the current noise:

$$\frac{\partial e_t^*(\theta)}{\partial \tau_{\epsilon t}} = D_t \alpha_t \sum_{i=t}^{T} \frac{b_i[\tau_{\theta i} - \alpha_t^2 \tau_{\epsilon t}]}{\tau_{\theta,i}^2} > 0,$$

where

$$D_t = \left((c_t')^{-1}\right)'\left(\alpha_t \tau_{\epsilon t} \sum_{i=t}^{T} \frac{b_i}{\tau_{\theta,i}}\right) = \frac{1}{c_t''\left((c_t')^{-1}\left(\alpha_t \tau_{\epsilon t} \sum_{i=t}^{T} \frac{b_i}{\tau_{\theta,i}}\right)\right)}.$$

Intuitively, the lower the variance of the noise, the more effort affects belief, so the greater the benefit of effort. Effort decreases in the precision of future noise.

Effort is independent of the mean of the belief, decreases in marginal cost[3], the precision of the belief and the future type differentiation $\alpha_{t+k},\ k \geq 1$. It is

[3] If the marginal cost function increases everywhere, then the effort of all types decreases.

intuitive that effort should decrease in marginal cost. If belief is more precise, then it is more difficult to change its mean, so the benefit of effort is lower. Based on (5), higher α_{t+k} reduces effort by increasing the precision of future beliefs, which lowers their responsiveness to effort.

The effect of the current type differentiation parameter α_t is

$$\frac{\partial e_t^*(\theta)}{\partial \alpha_t} = \theta + \tau_{\epsilon t} D_t \left(\sum_{i=t}^{T} \frac{b_i}{\tau_{\theta,i}} - 2\alpha_t^2 \tau_{\epsilon t} \sum_{i=t}^{T} \frac{b_i}{\tau_{\theta,i}^2} \right), \tag{7}$$

so effort increases in α_t for types above a cutoff and decreases for types below. The cutoff is generally nonzero, so some types $\theta < 0$ whose cost and marginal cost increase when α_t rises may exert more effort. For other parameter values, some types $\theta > 0$ may lower effort when α_t rises. If $|\theta|$ is large enough, the sign of $\frac{\partial e_t^*(\theta)}{\partial \alpha_t}$ is the same as the sign of θ, so effort rises for the highest and falls for the lowest types when the differentiation parameter α_t increases.

The change in effort as current type differentiation rises is driven by two forces. The first is the direct effect of a change in marginal cost (captured by θ on the RHS of (7)). This effect is positive iff $\theta > 0$. The second is the effect of the change in marginal benefit, which may be positive or negative. It is the product of the effect of current effort on belief at the end of the current period and the influence of this belief on the total benefit. As α_t rises, effort may influence belief at the end of period t more or less, depending on the balance of two forces: the greater difference of the expected effort across types (which raises effort) and the increased belief precision $\tau_{\theta,t}$ (which lowers effort). With an increase in α_t, the persistence of belief increases, so $\mu_{\theta,t}$ affects the total benefit more.

Increasing patience is described by the growth of δ in $b_n = \delta^{n-t}\hat{b}_n \ \forall n \geq t$. Only the current cost c_t matters for $e_t^*(\theta)$, so a change in the discount factor does not influence effort through cost.

$$\frac{\partial e_t^*(\theta)}{\partial \delta} = D_t \alpha_t \tau_{\epsilon t} \sum_{n=t+1}^{T} \frac{(n-t)\delta^{n-t-1}\hat{b}_n}{\tau_{\theta,t-1} + \sum_{k=t}^{n} \tau_{\epsilon k}\alpha_k^2} > 0.$$

Patience increases signalling effort, which is not surprising, because signalling is like investment—the cost is paid immediately, but the benefit is obtained in the future. Increased patience raises the weight the agent assigns to the future benefit in the payoff.

The value $V_{\theta,t}(\mu_{\theta,t-1}, \tau_{\theta,t-1})$ at the start of period t is the sum of expected benefits minus costs over future periods

$$\sum_{n=t}^{T} \left[b_n \frac{\tau_{\theta,t-1}\mu_{\theta,t-1} + \theta \sum_{k=t}^{n} \tau_{\epsilon k}\alpha_k^2}{\tau_{\theta,n}} - c_n \left((c_n')^{-1} \left(\alpha_n \tau_{\epsilon n} \sum_{i=n}^{T} \frac{b_i}{\tau_{\theta,i}} \right) \right) \right]. \tag{8}$$

The term $\frac{\tau_{\theta,t-1}\mu_{\theta,t-1} + \theta \sum_{k=t}^{n} \tau_{\epsilon k}\alpha_k^2}{\tau_{\theta,n}} =: \mathbb{E}_t \mu_{\theta,n}$ is the expected mean of the belief at the end of period n conditional on the information at the start of period

t. Value at t clearly increases in type and the mean of the belief (therefore also in the mean of the prior).

As can be seen from

$$\frac{\partial V_{\theta,t}(\mu_{\theta,t-1},\tau_{\theta,t-1})}{\partial b_n} = \frac{\tau_{\theta,t-1}\mu_{\theta,t-1} + \theta\sum_{k=t}^{n}\tau_{\epsilon k}\alpha_k^2}{\tau_{\theta,n}} - \frac{1}{\tau_{\theta,n}}\sum_{j=t}^{n}\tau_{\epsilon j}^2\alpha_j^2 D_j \sum_{i=j}^{T}\frac{b_i}{\tau_{\theta,i}}, \tag{9}$$

if the type or the mean of the belief is above a cutoff, then value increases in b_n, otherwise decreases. The first term on the RHS of (9) is the direct effect of a higher benefit on the value. The second term is the cost imposed by the increased effort motivated by higher marginal benefit.

Suppose $c_n(\cdot) = \eta\hat{c}_n(\cdot)$ with $\eta > 0$. The effect of cost on value is

$$\begin{aligned}&\frac{\partial V_{\theta,t}(\mu_{\theta,t-1},\tau_{\theta,t-1})}{\partial \eta}\\ &= -\hat{c}_n\left((\eta\hat{c}_n')^{-1}\left(\alpha_n\tau_{\epsilon n}\sum_{i=n}^{T}\frac{b_i}{\tau_{\theta,i}}\right)\right) + \frac{\alpha_n\tau_{\epsilon n}}{\eta^2}(\hat{c}_n')^{-1}\left(\alpha_n\tau_{\epsilon n}\sum_{i=n}^{T}\frac{b_i}{\tau_{\theta,i}}\right)\sum_{j=n}^{T}\frac{b_j}{\tau_{\theta,j}}.\end{aligned}$$

This may be positive or negative, depending on the balance between the direct effect of increased cost (the first term on the RHS) and the indirect effect of lower effort (the second term).

Write the cost functions and benefit parameters again as $c_n(\cdot) = \delta^{n-t}\hat{c}_n$, $b_n = \delta^{n-t}\hat{b}_n$ $\forall n \geq t$. The effect of increased patience (higher δ) on the value is

$$\begin{aligned}\frac{\partial V_{\theta,t}(\mu_{\theta,t-1},\tau_{\theta,t-1})}{\partial \delta} &= \sum_{n=t+1}^{T}(n-t)\delta^{n-t-1}\hat{b}_n\frac{\tau_{\theta,t-1}\mu_{\theta,t-1} + \theta\sum_{k=t}^{n}\tau_{\epsilon k}\alpha_k^2}{\tau_{\theta,n}}\\ &- \sum_{n=t+1}^{T}(n-t)\delta^{n-t-1}\hat{c}_n\left((\hat{c}_n')^{-1}\left(\alpha_n\tau_{\epsilon n}\sum_{i=n}^{T}\frac{\delta^{i-t}\hat{b}_i}{\tau_{\theta,i}}\right)\right)\\ &- \sum_{n=t}^{T}\delta^{n-t}\alpha_n\tau_{\epsilon n}\sum_{i=n}^{T}\frac{\delta^{i-n}\hat{b}_i}{\tau_{\theta,i}}\left[\hat{c}_n''\left((\hat{c}_n')^{-1}\left(\alpha_n\tau_{\epsilon n}\sum_{j=n}^{T}\frac{\delta^{j-n}\hat{b}_j}{\tau_{\theta,j}}\right)\right)\right]^{-1}\\ &\times \alpha_n\tau_{\epsilon n}\sum_{k=n+1}^{T}\frac{(k-n)\delta^{k-n-1}\hat{b}_k}{\tau_{\theta,k}}.\end{aligned}$$

If the type or the mean of the belief is above a cutoff, then value increases in δ, otherwise decreases. The first two sums on the RHS describe the increase in the discounted expected benefit and cost as the discount rate increases. The last sum is the effect of a change in effort in response to an altered discount rate.

The reason the value of the highest types increases in δ and the value of the lowest decreases is that a rise in the discount factor puts more weight on the future payoffs relative to the present. Low types expect their stage game payoff to decrease in the future and high types expect it to increase.

The intuition why value increases in δ for high beliefs is that the flow payoff has the same sign as the mean belief and belief is persistent. For high beliefs, positive flow benefit is expected in the near future. This benefit has more weight with a higher δ. Note that value has not been divided by $1-\delta$ to control for the increase in the absolute value of the sum as the elements get multiplied by an increasing δ. Defining $\hat{V} = \frac{V}{1-\delta}$, the derivative is $\frac{\partial \hat{V}_{\theta,t}}{\partial \delta} = \frac{1}{1-\delta}\frac{\partial V_{\theta,t}}{\partial \delta} + \frac{1}{(1-\delta)^2}V_{\theta,t}$. Similarly to $\frac{\partial V_{\theta,t}}{\partial \delta}$, this is positive if the type or the mean of the belief is above a cutoff.

The effect of increased precision of period-j noise on the value is

$$\frac{\partial V_{\theta,t}}{\partial \tau_{\epsilon j}} = \sum_{n=t}^{T} \frac{b_n \tau_\theta (\theta - \mu_\theta)\alpha_j^2}{\tau_\theta + \sum_{k=1}^{n} \tau_{\epsilon k}\alpha_k^2} - \sum_{n=t}^{T} \alpha_n^2 \tau_{\epsilon n} \sum_{i=n}^{T} \frac{b_i}{\tau_{\theta,i}} D_n \sum_{i=n}^{T} \frac{b_i(\mathbf{1}_{j=n}\tau_{\theta,i} - \alpha_j^2 \tau_{\epsilon n})}{\tau_{\theta,i}^2},$$

which is negative if $\theta < \mu_{\theta,t-1}$ and unclear otherwise. More precise noise speeds convergence of belief to the true type and increases effort, which increases cost. For $\theta < \mu_{\theta,t-1}$, both these effects are negative, but for $\theta > \mu_{\theta,t-1}$, convergence to the true type increases payoff.

The derivative of the value w.r.t. the precision of the belief is

$$\frac{\partial V_{\theta,t}}{\partial \tau_{\theta,t-1}} = \sum_{n=t}^{T} b_n \frac{(\mu_{\theta,t-1} - \theta)\sum_{k=t}^{n} \tau_{\epsilon k}\alpha_k^2}{\left(\tau_{\theta,t-1} + \sum_{k=t}^{n} \tau_{\epsilon k}\alpha_k^2\right)^2} + \sum_{n=t}^{T} \alpha_n \tau_{\epsilon n} \sum_{i=n}^{T} \frac{b_i}{\tau_{\theta,i}} D_n \alpha_n \tau_{\epsilon n} \sum_{k=n}^{T} \frac{b_k}{\tau_{\theta,k}^2}$$

This is positive if the mean of the belief is greater than the type, and has unclear sign otherwise. A more precise belief responds less to future signals (is more persistent), so an overestimation of the type by the market is corrected slower. This is captured by the first term on the RHS. The second term describes the effect on the value of a reduction in effort in response to a more precise belief. Lower effort decreases cost, thereby increasing the value.

The effect of the period-j cost difference $\alpha_j,\ j \geq t$ between the types on the value is

$$\frac{\partial V_{\theta,t}}{\partial \alpha_j} = 2\alpha_j \tau_{\epsilon j}\tau_{\theta,t-1} \sum_{n=j}^{T} b_n \frac{\theta - \mu_{\theta,t-1}}{\tau_{\theta,n}^2} - \alpha_j \tau_{\epsilon j}^2 D_j \left(\sum_{k=j}^{T} \frac{b_k}{\tau_{\theta,k}}\right)^2$$
$$+ 2\alpha_j \tau_{\epsilon j} \sum_{n=t}^{T} \alpha_n^2 \tau_{\epsilon n}^2 \sum_{i=n}^{T} \frac{b_i}{\tau_{\theta,i}} \left[c_n'' \left((c_n')^{-1}\left(\alpha_n \tau_{\epsilon n} \sum_{l=n}^{T} \frac{b_l}{\tau_{\theta,l}}\right)\right)\right]^{-1} \sum_{k=j}^{T} \frac{b_k}{\tau_{\theta,k}^2}.$$

The response of the value to a higher α_j is positive if $\theta - \mu_{\theta,t-1}$ is above a cutoff and negative otherwise. Depending on the parameters, the value may go up for some types $\theta < 0$ whose cost goes up when α_j increases or go down when cost falls.

5 Extensions

5.1 Exogenously changing type

Employees could have good and bad days, which make their ability change over time. This can be modelled as an exogenously varying type. Similarly, a politician's competence in the foremost issue of the day varies exogenously if the issue in question changes independently of the politician, e.g. a crisis arises due to the actions of one neighbouring country or another.

Assume type changes as a random walk $\theta_t = \theta_{t-1} + \nu_t$, with $\nu_t \sim N(0, \tau_\nu)$ independently of everything else. The distribution of θ_0 is $N(\mu_{\theta,0}, \tau_{\theta,0})$. Assume the sender observes θ_{t-1} when choosing e_t. Let θ_t denote the type at the end of period t.

The market update their belief about e_t based on s_t the same way as before. The updating formula for θ_{t-1} based on s_t is still (4), where θ has acquired the subscript $t-1$ and γ_{T-t} has changed. Since the mean of the belief as a function of effort is the same as in the constant type case, the marginal flow benefit of effort is given by the same formula: $m_t = \frac{b_t \alpha_t \tau_{\epsilon t}}{\tau_{\theta,t-1} + \alpha_t^2 \tau_{\epsilon t}}$. From the perspective of the market at the end of period t, the mean of θ_t is the same as for θ_{t-1}, but the precision is

$$\tau_{\theta,t} = \frac{\tau_\nu(\tau_{\theta,t-1} + \alpha_t^2 \tau_{\epsilon t})}{\tau_\nu + \tau_{\theta,t-1} + \alpha_t^2 \tau_{\epsilon t}}.$$

The marginal benefit of signalling in period t can still be decomposed into the part m_t received in the current period and the part obtained in the future. The latter is the product of three components. The first is $\frac{\alpha_t \tau_{\epsilon t}}{\tau_{\theta,t-1} + \alpha_t^2 \tau_{\epsilon t}}$, the influence of the sender's current effort on the mean of the market's belief at the end of the current period. The second is $\prod_{k=1}^{n-1} \frac{\tau_{\theta,t+k-1}}{\tau_{\theta,t+k-1} + \alpha_{t+k}^2 \tau_{\epsilon,t+k}}$, the effect of the mean of the belief at the end of period t on the mean of the belief at the end of period $t+n-1$ (this component is absent for $n = 1$). The third is $\frac{b_{t+n} \tau_{\theta,t+n-1}}{\tau_{\theta,t+n-1} + \alpha_{t+n}^2 \tau_{\epsilon,t+n}}$, the change in period $t+n$ benefit in response to a change in the mean of the belief at the end of period $t+n-1$. The total future marginal benefit of effort is

$$MB_t^{fut} = \frac{\alpha_t \tau_{\epsilon t}}{\tau_{\theta,t-1} + \alpha_t^2 \tau_{\epsilon t}} \sum_{n=1}^{T} b_{t+n} \prod_{k=1}^{n} \frac{\tau_{\theta,t+k-1}}{\tau_{\theta,t+k-1} + \alpha_{t+k}^2 \tau_{\epsilon,t+k}}.$$

The new γ_{T-t} is $(c_t')^{-1}\left(m_t + MB_t^{fut}\right)$.

The optimal strategy is

$$e_t^*(\theta_{t-1}) = (c_t')^{-1}\left(\frac{\alpha_t \tau_{\epsilon t}}{\tau_{\theta,t-1} + \alpha_t^2 \tau_{\epsilon t}} \sum_{n=0}^{T} b_{t+n} \prod_{k=1}^{n} \frac{\tau_{\theta,t+k-1}}{\tau_{\theta,t+k-1} + \alpha_{t+k}^2 \tau_{\epsilon,t+k}}\right) + \alpha_t \theta_{t-1},$$

so the only difference from (5) is in $\tau_{\theta,n}$.

In the value $V_{\theta,t}(\mu_{\theta,t-1}, \tau_{\theta,t-1})$ at the start of period t, the expected future type equals the current type, $\mathbb{E}_t\theta_{t+n} = \theta_{t-1}$. The cost at the equilibrium effort is the same for all types. The formula for the value is

$$\sum_{n=t}^{T}\left[b_n\frac{\tau_{\theta,n-1}\mu_{\theta,t-1}+\alpha_n^2\tau_{\epsilon n}\theta_{t-1}}{\tau_{\theta,n-1}+\alpha_n^2\tau_{\epsilon n}} - c_n\left(\gamma_{T-t}\right)\right].$$

5.1.1 Stationary environment

Take $T=\infty$, $\alpha_t=\alpha$ (constant over time), $b_t=\delta^t$ (divide both cost and benefit by b to normalize $b=1$) and $c_t(\cdot)=\delta^t c(\cdot)$, with $\delta\in(0,1)$. W.l.o.g. take $\tau_\epsilon=1$ (divide τ_θ by τ_ϵ). The relation

$$\tau_{\theta,0}=\frac{\tau_\nu(\tau_{\theta,0}+\alpha^2)}{\tau_\nu+\tau_{\theta,0}+\alpha^2}$$

between the initial precision $\tau_{\theta,0}$ of the type distribution and the precision τ_ν of the random type change ensures that $\tau_{\theta,t}$ is constant over time at the level $\tau_\theta=-\frac{\alpha^2}{2}+\sqrt{\frac{\alpha^4}{4}+\tau_\nu\alpha^2}$.

The future marginal benefit of effort is

$$MB_t^{fut}=\frac{\alpha\delta\tau_\theta}{(\tau_\theta+\alpha^2)[(1-\delta)\tau_\theta+\alpha^2]}.$$

The optimal effort in period t is

$$e^*(\theta_{t-1})=(c')^{-1}\left(\frac{\alpha}{(1-\delta)\tau_\theta+\alpha^2}\right)+\alpha\theta_{t-1}.$$

Effort increases in θ_{t-1} and δ, and decreases in marginal cost and τ_θ. Effort increases in α iff the type is above a cutoff. The steady state precision τ_θ of the type distribution rises in the precision τ_ν of the type innovations. Due to this, effort falls in τ_ν.

The value from period t on is

$$\frac{\tau_\theta\mu_{\theta,t-1}+\alpha^2\theta_{t-1}}{(1-\delta)(\tau_\theta+\alpha^2)}-\frac{1}{1-\delta}c\left((c')^{-1}\left(\frac{\alpha}{(1-\delta)\tau_\theta+\alpha^2}\right)\right).$$

The value rises in $\mu_{\theta,t-1}$ and θ_{t-1}. Adding a positive constant to the cost function c reduces value. The effects of other parameters are less clear, because they affect both cost and benefit. If θ_{t-1} or $\mu_{\theta,t-1}$ is large enough, then $\frac{\partial V}{\partial\delta}$ is positive, otherwise negative. The derivatives $\frac{\partial V}{\partial\alpha}$, $\frac{\partial V}{\partial\tau_\theta}$ and $\frac{\partial V}{\partial\tau_\nu}$ are positive if $\theta_{t-1}-\mu_{\theta,t-1}$ is large enough and negative otherwise. Multiplying the cost function by $\eta>1$ may increase or decrease value. All these comparative statics accord with those in the nonstationary environment.

5.2 Human capital accumulation

Employers sometimes care about a combination of the employee's intelligence and qualifications. This summary measure of employee quality rises with higher effort to acquire education. Investment in education thus has two purposes—to signal one's type to the employers and to improve that type.

Suppose that the type next period depends on current type and effort according to $\theta_t = \theta_{t-1} + h_t e_t$, with $h_t > 0$ and $\sum_{k=0}^{T} h_k < \infty$. The fact that effort can be negative neatly captures the depreciation of human capital if too little effort is exerted to maintain it (type falls with negative effort).

Conjecture that equilibrium effort is affine in type: $e_t(\theta_{t-1}) = \gamma_{T-t} + \alpha_t \theta_{t-1}$. Given that the prior on θ_{t-1} is normal and that in equilibrium, type evolves according to $\theta_t = (1+\alpha_t h_t)\theta_{t-1} + h_t\gamma_{T-t}$, the belief about θ_t *before* receiving signal s_t is

$$\theta_t | s_{t-1} \sim N\left((1+\alpha_t h_t)\mu_{\theta,t-1} + h_t\gamma_{T-t}, \frac{\tau_{\theta,t-1}}{(1+\alpha_t h_t)^2}\right).$$

Expressing $\theta_{t-1} = \frac{\theta_t - h_t\gamma_{T-t}}{1+\alpha_t h_t}$, effort can be written as $e_t = \gamma_{T-t} + \frac{\alpha_t(\theta_t - h_t\gamma_{T-t})}{1+\alpha_t h_t}$, so the signal distribution conditional on θ_t is normal with precision $\tau_{\epsilon t}$ and mean $\gamma_{T-t} + \frac{\alpha_t(\theta_t - h_t\gamma_{T-t})}{1+\alpha_t h_t} = \frac{\alpha_t\theta_t}{1+\alpha_t h_t} + \frac{\gamma_{T-t}}{1+\alpha_t h_t}$. Linearly transform the signal s_t to $z_t = \frac{(1+\alpha_t h_t)s_t - \gamma_{T-t}}{\alpha_t}$, which has mean θ_t and precision $\frac{\alpha_t^2\tau_{\epsilon t}}{(1+\alpha_t h_t)^2}$. Then the updated type at the end of period t is

$$\theta_t | z_t \sim N\left(\frac{\tau_{\theta,t-1}\left[(1+\alpha_t h_t)\mu_{\theta,t-1} + h_t\gamma_{T-t}\right] + \alpha_t^2\tau_{\epsilon t} z_t}{\tau_{\theta,t-1} + \alpha_t^2\tau_{\epsilon t}}, \frac{\tau_{\theta,t-1} + \alpha_t^2\tau_{\epsilon t}}{(1+\alpha_t h_t)^2}\right).$$

The precision of the belief at the end of period t is $\tau_{\theta,t} = \frac{\tau_{\theta,0}}{\prod_{n=1}^{t}(1+\alpha_n h_n)^2} + \sum_{n=1}^{t} \frac{\alpha_n^2\tau_{\epsilon n}}{\prod_{k=n}^{t}(1+\alpha_k h_k)^2}$, smaller than without human capital accumulation. Since higher types accumulate human capital faster, the type distribution becomes more dispersed over time. This counteracts learning by the market and may even make the precision of the posterior belief decrease in time.

Given effort e_t, the expected benefit in period t is

$$b_t \frac{\tau_{\theta,t-1}\left[(1+\alpha_t h_t)\mu_{\theta,t-1} + h_t\gamma_{T-t}\right] + \alpha_t\tau_{\epsilon t}[(1+\alpha_t h_t)e_t - \gamma_{T-t}]}{\tau_{\theta,t-1} + \alpha_t^2\tau_{\epsilon t}},$$

so the marginal flow benefit of effort is $\frac{b_t\alpha_t\tau_{\epsilon t}(1+\alpha_t h_t)}{\tau_{\theta,t-1}+\alpha_t^2\tau_{\epsilon t}}$. The marginal benefit at time t of shifting μ_{t-1} is $\frac{b_t\tau_{\theta,t-1}(1+\alpha_t h_t)}{\tau_{\theta,t-1}+\alpha_t^2\tau_{\epsilon t}}$. The response of μ_{t+n-1} ($n \geq 2$) to μ_t is

$$\frac{\partial\mu_{t+n-1}}{\partial\mu_t} = \frac{\tau_{\theta,t}}{\tau_{\theta,t+n}} \prod_{k=0}^{n-1} \frac{1}{1+\alpha_{t+k+1}h_{t+k+1}}.$$

The effect of e_t on μ_t is $\frac{\alpha_t\tau_{\epsilon t}(1+\alpha_t h_t)}{\tau_{\theta,t-1}+\alpha_t^2\tau_{\epsilon t}}$. The total marginal benefit of effort is

$$MB_t = \alpha_t\tau_{\epsilon t}\sum_{n=0}^{T-t}\frac{b_{t+n}}{\frac{\tau_{\theta,0}}{\prod_{j=1}^{t-1}(1+\alpha_j h_j)^2\prod_{j=t}^{t+n}(1+\alpha_j h_j)}+\sum_{j=1}^{t+n}\frac{\alpha_j^2\tau_{\epsilon j}}{\prod_{k=j}^{t-1}(1+\alpha_k h_k)^2\prod_{k=t}^{t+n}(1+\alpha_k h_k)}}.$$

This is constant in type and effort, verifying the conjecture that equilibrium effort $e_t^*(\theta_{t-1}) = (c_t')^{-1}(MB_t)+\alpha_t\theta_{t-1}$is affine in type. Effort at time t increases in h_j for all j, which is intuitive, because human capital accumulation provides an extra benefit of effort.

The effect of faster human capital accumulation on the value is unclear, because both cost and benefit increase in h_j.

5.3 Exogenous information revelation

The skill of an employee may be revealed not just through education signalling, but also by word of mouth. The profit prospects of a firm can be published in a report by a rating agency or stock analyst, not just signalled via taking on debt or repurchasing shares. The quality of a product can be inferred from independent reviews as well as from ads or warranties used as signalling devices. To describe these situations, an exogenous signal $x_t = \theta + \xi_t$ is added to the baseline model, with $\xi_t \sim N(0, \tau_\xi)$ independent of all other variables. Assume the benefit is obtained at the end of each period after both s_t, x_t have been observed. Without loss of generality, assume x_t occurs after the endogenous signal s_t in each period t.

Denote by $\mu_{\theta,t}^s$ and $\tau_{\theta,t}^s$ the mean and precision of the belief after observing s_t, but before x_t. The mean and precision after observing x_t are written $\mu_{\theta,t}, \tau_{\theta,t}$. Belief is updated based on s_t by (4), so the belief at the end of period t is

$$\begin{aligned}&N\left(\frac{\mu_{\theta,t}^s\tau_{\theta,t}^s + x_t\tau_\xi}{\tau_{\theta,t}^s+\tau_\xi},\quad \tau_{\theta,t}^s+\tau_\xi\right)\\ &= N\left(\frac{\tau_{\theta,t-1}\mu_{\theta,t-1}+\alpha_t\tau_{\epsilon t}(s_t-\gamma_{T-t})+x_t\tau_\xi}{\tau_{\theta,t-1}+\alpha_t^2\tau_{\epsilon t}+\tau_\xi},\quad \tau_{\theta,t-1}+\alpha_t^2\tau_{\epsilon t}+\tau_\xi\right).\end{aligned}$$

The calculations for the optimal effort are the same as in the baseline model. The only difference is that $\tau_{\theta,t} = \tau_\theta + t\tau_\xi + \sum_{k=1}^t \tau_{\epsilon k}\alpha_k^2$. The comparative statics of effort are thus similar to the baseline, with the departure that in all expressions, $\tau_{\theta,t}$ is larger for each t. A higher precision τ_ξ of the exogenous signal reduces effort, because it increases the precision of the belief in all periods. Intuitively, free information revelation reduces the incentive to convey the information via costly signalling.

The value from period t on is

$$\sum_{k=0}^{T-t} b_{t+k} \frac{\tau_{\theta,t-1}\mu_{\theta,t-1} + \left[(k+1)\tau_\xi + \sum_{i=t}^{t+k} \tau_{\epsilon i}\alpha_i^2\right]\theta}{\tau_{\theta,t-1} + (k+1)\tau_\xi + \sum_{i=t}^{t+k} \tau_{\epsilon i}\alpha_i^2}$$
$$- \sum_{k=0}^{T-t} c_{t+k}\left((c'_{t+k})^{-1}\left(\alpha_{t+k}\tau_{\epsilon,t+k}\sum_{n=t+k}^{T}\frac{b_n}{\tau_{\theta,n}}\right)\right).$$

This expression is similar to (8)—the only changes reflect the faster learning using the exogenous signal. The comparative statics w.r.t. the variables previously present are thus similar. The only new effect is that of τ_ξ. Value increases in τ_ξ iff $\theta - \mu_{\theta,t-1}$ is above a cutoff. The cutoff is always positive, so a sufficient condition for value to rise in τ_ξ is $\theta > \mu_{\theta,t-1}$. If the type is above the belief, then faster information revelation raises the benefit quicker. In addition, more exogenous type revelation reduces signalling effort, which decreases cost.

5.3.1 One signal depending on both type and effort

Grades may depend on both intelligence and effort, and these two components may not be distinguishable based on the diploma or transcript. To model this, suppose the signal depends on type according to $s_t = r\theta + e_t + \epsilon_t$, with $r \in \mathbb{R}_+$. If an affine strategy $e_t^*(\theta) = k_{1t} + k_{2t}\theta$ is expected, with $k_{2t} > 0\ \forall t$, then the updating formula is

$$\theta|s_t \sim N\left(\frac{\tau_{\theta,t-1}\mu_{\theta,t-1} + (r+k_{2t})\tau_{\epsilon t}(s_t - k_{1t})}{\tau_{\theta,t-1} + (r+k_{2t})^2\tau_{\epsilon t}},\quad \tau_{\theta,t-1} + (r+k_{2t})^2\tau_{\epsilon t}\right).$$

The marginal benefit calculation is similar to the basic model, with $\tau_{\theta,t} = \tau_\theta + \sum_{i=1}^{t}\tau_{\epsilon i}(r+\alpha_i)^2$. Equilibrium effort is

$$e_t^*(\theta) = (c'_t)^{-1}\left((r+\alpha_t)\tau_{\epsilon t}\sum_{n=t}^{T}\frac{b_n}{\tau_{\theta,t-1} + \sum_{k=t}^{n}\tau_{\epsilon k}(r+\alpha_k)^2}\right) + \alpha_t\theta.$$

Effort may increase or decrease in r. On the one hand, a higher r makes future beliefs more precise (less responsive to effort), but on the other hand, makes belief at the end of the current period change more in effort.

The value from period t on is similar to (8), with the exception of α_n being replaced by $r + \alpha_n$. The comparative statics in r are the same as changing all α_t-s simultaneously in (8). As r rises, both cost and benefit may increase or decrease. The cost effect operates via effort. On the benefit side, a higher r raises the speed of learning, which increases payoff iff $\theta > \mu_{\theta,t-1}$.

5.4 Productive effort

The effort that workers use to signal their talent can be directly productive to the employer. For example, effort at work could have a signalling component, with the employer using it to learn about the quality of the worker. This situation is similar to that examined in the career concerns literature, e.g. Holmström (1999), who assumes type is unknown to both worker and employer, signal depends on both type and effort and the employer values both type and effort. In this section, it is assumed that type is known to the worker and the signal depends only on effort.

Competition between risk-neutral employers will drive the wage of the worker to $w = \mathbb{E}\left[q_t\theta + (1-q_t)e_t^*(\theta)\right]$ in any period t, where $q_t \in (0,1)$ is the relative weight of type in productivity. The worker is risk-neutral and gets utility $b_t w$ from w.

Conjecture an affine strategy $e_t^*(\theta) = k_{1t} + k_{2t}\theta$, with $k_{2t} > 0\ \forall t$. Updating based on the signal $s_t = e_t + \epsilon_t$ yields

$$\theta|s_t \sim N\left(\frac{\tau_{\theta,t-1}\mu_{\theta,t-1} + k_{2t}\tau_{\epsilon t}(s_t - k_{1t})}{\tau_{\theta,t-1} + k_{2t}^2\tau_{\epsilon t}}, \quad \tau_{\theta,t-1} + k_{2t}^2\tau_{\epsilon t}\right).$$

The expected wage given effort e is then

$$(1-q_t)k_{1t} + (q_t + (1-q_t)k_{2t})\frac{\tau_{\theta,t-1}\mu_{\theta,t-1} + k_{2t}\tau_{\epsilon t}(e - k_{1t})}{\tau_{\theta,t-1} + k_{2t}^2\tau_{\epsilon t}},$$

because the expected signal equals the effort. The marginal flow benefit of effort is $\frac{b_t(q_t+(1-q_t)k_{2t})k_{2t}\tau_{\epsilon t}}{\tau_{\theta,t-1}+k_{2t}^2\tau_{\epsilon t}}$. Current effort shifts subsequent beliefs as well and this provides the marginal future benefit $\frac{k_{2t}\tau_{\epsilon t}}{\tau_{\theta,t-1}+k_{2t}^2\tau_{\epsilon t}}\sum_{n=t+1}^{T}\frac{b_n(q_n+(1-q_n)k_{2n})\tau_{\theta,t}}{\tau_{\theta,n}}$. The total marginal benefit is constant, so the equilibrium effort is

$$e_t^*(\theta) = (c_t')^{-1}\left(\alpha_t\tau_{\epsilon t}\sum_{n=t}^{T}\frac{b_n(q_n+(1-q_n)\alpha_n)}{\tau_{\theta,t-1}+\sum_{k=t}^{n}\tau_{\epsilon k}\alpha_k^2}\right) + \alpha_t\theta,$$

affine in type as required, with $k_{2t} = \alpha_t\ \forall t$ as in the basic model. The difference from (5) is only the constants $q_n + (1-q_n)\alpha_n$.

Raising any q_n, $n \geq t$ raises $e_t^*(\theta)$ iff $\alpha_n < 1$. In words, increasing the weight of type in the wage raises effort if the difference between the equilibrium efforts of the types is less than the type difference. Based on the wage expression, type is productive directly and via effort. If $\alpha_t < 1$, then in period t, the difference in the direct productivity of the types is greater than the gap in their efforts. Increasing the weight q_t of this direct effect in wage then increases the responsiveness of wage to type and thereby the incentive to exert effort.

Value increases in q_t iff a convex combination of $\mu_{\theta,t}$ and θ is above a cutoff. The cost component of value responds to q_t via equilibrium effort. Thus if $\alpha_t < 1$, then a higher q_t tends to raise cost and reduce value. The benefit component contains a type-dependent and a type-independent part. The type-dependent part increases in q_t iff $\alpha_t < 1$, because then the direct effect of type on productivity is greater than the indirect effect through effort. Raising q_t increases the relative weight of the direct effect. The type-independent part of benefit relies on effort. Increasing q_t reduces the relative size of this part, but also changes effort. The latter effect raises the type-independent part iff $\alpha_t < 1$.

5.5 Multiple signallers

Researchers signal their talent by publishing papers or obtaining patents, both of which are stochastically increasing in effort. If other researchers produce more results in the same field, then it becomes more difficult to publish work of a given quality. High signals of others thus increase one's cost of signalling.

Work in a complementary field may make research effort less costly in a given discipline. New microscopes, DNA sequencing machines or statistical techniques all facilitate discovery. This can be modelled as high signals of others reducing the cost of signalling.

There are two senders A and B, with cost $c\left(e_{it} - \alpha\theta_i - \sum_{k=1}^{t-1} \kappa_{ikt} s_{jk}\right)$ for sender i influenced by the past signals s_{jk} of sender $j \neq i$. For simplicity, b, c and α are constant over time and across senders. The influence κ_{ikt} of the other sender's success in period k on i's cost in t may be positive or negative.

The solution procedure is the same as in the previous sections. The equilibrium effort is

$$e_{it}^{*}(\theta_i) = (c')^{-1}\left(\sum_{n=t}^{T} \frac{\alpha\tau_{\epsilon}b}{\tau_{i\theta,t-1} + \tau_{\epsilon}(n-t+1)\alpha^2}\right) + \alpha\theta_i + \sum_{k=1}^{t-1} \kappa_{ikt} s_{jk}.$$

It would seem *a priori* that each player wants to manipulate the effort of the other to obtain lower cost for himself in the future. This is not true for the particular form of cost interaction considered here. Modifying the cost function of a player from $c(\cdot)$ to $c(\cdot + k)$, $k \in \mathbb{R}$, changes the equilibrium effort of all types by $-k$. The cost paid in equilibrium is unchanged: $c(e^{*}(\theta) - k + k) = c(e^{*}(\theta))$. The altered effort shifts the signal distribution, which would be expected to change the benefit. However, the signal distributions of all types shift by the same constant and retain their original shape, so Bayes' rule leads to the same belief distribution as before. The benefit (which depends on the belief, not the signal directly) is therefore unchanged. Overall, neither value nor incentives respond to a cost change from $c(\cdot)$ to $c(\cdot + k)$, so each player is indifferent about influencing the cost of the other player or their own future self this way.

If a player's effort or signal in some period affects the other player's same-period cost, then the model is not tractable. This is because an expectation of the (nonlinear) cost appears in the FOC, making the best response nonlinear in type. Bayesian updating then yields a complicated integral equation.

If the benefit for player i in period t is $b\mu_{i\theta t} + f_i(\mu_{j\theta t})$ for some function $f_i : \mathbb{R} \to \mathbb{R}$, then the optimal strategy is given by (5) as in the basic model. The additive effect of the belief about the other player disappears when taking the FOC and thus does not affect incentives.

5.6 Observed effort and two-dimensional private information

The central bank's choice of interest rate e_t is perfectly observed by the market in each period t. The central bank privately knows its 'ideal interest rate' θ

and tries to signal that this is high by setting a high interest rate.[4] The second dimension of private information that the central bank has is an i.i.d. signal $\epsilon_t \sim N(0, \tau_\epsilon)$ about the state of the economy. The better the state of the economy, the higher the interest rate the central bank would like to set.

There is a cost to setting a higher interest rate—this dampens economic activity in the current period. The benefit from raising e_t comes from the higher mean of the belief of the market about the type of the central bank, because higher belief reduces inflation expectations. The stage game payoff of the central bank is $b\mu_{\theta t} - c(e_t - \alpha\theta - \epsilon_t)$, where $c(\cdot)$ satisfies the same assumptions as in the baseline model.

If the marginal benefit MB_t in period t is constant in type and effort, then a central bank of type θ with the private signal ϵ has the best response $e_t(\theta, \epsilon) = (c')^{-1}(MB_t) + \alpha\theta + \epsilon_t$. If the market expects strategy $e_t^*(\theta, \epsilon) = k_{1t} + k_{2t}\theta + k_{3t}\epsilon_t$, then the updated type after interest rate e is

$$\theta | e_t \sim N\left(\frac{\tau_{\theta,t-1}\mu_{\theta,t-1}k_{3t}^2 + k_{2t}\tau_\epsilon(e_t - k_{1t})}{\tau_{\theta,t-1}k_{3t}^2 + k_{2t}^2\tau_\epsilon}, \quad \tau_{\theta,t-1} + \frac{k_{2t}^2}{k_{3t}^2}\tau_\epsilon\right).$$

The marginal flow benefit of effort is then $b\frac{k_{2t}\tau_\epsilon}{\tau_{\theta,t-1}k_{3t}^2 + k_{2t}^2\tau_\epsilon}$. The marginal benefit from shifting future belief is $\frac{k_{2t}\tau_\epsilon}{k_{3t}^2}\sum_{n=t+1}^{T}\frac{b}{\tau_{\theta n}}$, constant in type and effort as required. The equilibrium strategy is

$$e_t^*(\theta, \epsilon) = (c')^{-1}\left(\alpha\tau_\epsilon\sum_{n=t}^{T}\frac{b}{\tau_{\theta n}}\right) + \alpha\theta + \epsilon_t.$$

If the learning process starts anew every time the governor of the central bank changes, then the model predicts interest rates to be highest at the start of a governor's tenure, because $\tau_{\theta n}$ rises over time.

A firm that signals by observably taking on debt as in Leland and Pyle (1977) also has private information with some dimensions persistent (soundness of the core business) and some changeable (current day's demand). Extending Leland and Pyle (1977) to multiple periods using the current paper's framework, the capital structure is predicted to be furthest from optimal for the first equity issue. The debt level converges to optimal as the market learns the type of the firm.

Applying the model to limit pricing as in Milgrom and Roberts (1982), the sender's observable choice is the price, the persistent private information is the marginal cost and the time-varying private signal is the cash-flow (the need of the firm for immediate revenue). A newly established monopoly is predicted to set the lowest price to deter entry. The price will rise towards the monopoly level over time as the potential entrants learn of the cost of the incumbent.

With observed effort, a continuum of equilibria appears, as is common in noiseless signalling games. Belief threats attached to actions off the equilibrium path prevent any type from taking these actions. An arbitrarily small amount of noise in the action destroys these new equilibria.

[4] This story is based on Vickers (1986).

6 Literature

A large literature on signalling started from the seminal paper of Spence (1973). The relevance of repetition in signalling was discussed in Weiss (1983) and Admati and Perry (1987) already. The importance of noise was made clear in Matthews and Mirman (1983). The present paper is a natural combination of these ideas, as are several previous publications.

The works of Kaya (2009) and Roddie (2012) examine discrete time repeated signalling without noise. These papers are distinguished from the present paper by the perfect observability of the sender's action. The focus in Kaya (2009) is on least-cost separating equilibria. The paper of Roddie (2012) shows when reputation effects arise in general. In the noiseless models of Kaya (2009) and Roddie (2012), the existence and payoffs of separating equilibria are independent of the prior, and pooling on positive effort is possible. These results differ from the noisy models discussed below, as well as the present paper.

Discrete time noisy signalling is discussed in Dilme and Li (2014). In that work, the decision is when to irreversibly stop signalling. Noise is one-sided: a choice to continue can randomly result in stopping, but not vice versa. In the present paper, noise is two-sided: a given effort level can result in a higher or lower signal. The decision is how much effort to exert each period, not when to stop.

In the discrete time model of Dilme (2014a), the high type prefers high effort even in the absence of signalling considerations. The low type chooses when to stop imitating the high. The focus is on characterizing the set of equilibrium payoffs. In the current paper, both types prefer lower effort if there is no signalling motive. Effort is chosen from a continuum and the game never stops. The goal is to describe behaviour in a particular equilibrium.

Continuous time signalling with Brownian noise is the subject of Daley and Green (2012), Gryglewicz (2009), Dilme (2014b) and Heinsalu (2014). In the paper of Dilme (2014b), the sender continuously chooses the effort level and the probability of stopping the game. The benefit is obtained only upon stopping. In the present paper, the benefit may be obtained in multiple periods or in one, and the benefit timing is known in advance. In Daley and Green (2012), information is exogenously revealed over time and the informed player decides when to stop the game. Again, the benefit is received upon stopping. Limit pricing in a dynamic environment is studied in Gryglewicz (2009). The low-cost type is nonstrategic and the high-cost type chooses the time at which it ceases its costly imitation of the low-cost type. Heinsalu (2014) has a model with two types and either a Poisson or a Brownian signal process. Effort is chosen continuously and determines the distribution of signals. In the current paper, all types are strategic, there is a continuum of types and the signal distribution depends on effort.

Repeated noiseless signalling is also studied in Nöldeke and van Damme (1990) and Swinkels (1999). The framework is similar to Dilme and Li (2014)—the signalling cost is paid first and the benefit is received at the end of the game. This feature differentiates them from the current paper, where benefit may be received in all periods. A unique informative equilibrium obtains in Nöldeke and van Damme (1990). The different information structure in Swinkels (1999) leads to a unique pooling equilibrium. The current paper has one affine equilibrium, but the existence of nonlinear equilibria is an open question.

Signalling by delaying trade is the subject of Kremer and Skrzypacz (2007) and Hörner and Vieille (2009). The benefit is obtained when trade occurs, i.e. when the game stops. In the present paper, there is no irreversible trading choice. Instead, there is a choice of effort every period.

Signalling in one-shot interactions is studied in Matthews and Mirman (1983), Carlsson and Dasgupta (1997), Alós-Ferrer and Prat (2012) and Daley and Green (2014). The present paper focusses on long-term relationships. Matthews and Mirman (1983) and Daley and Green (2014) add noise to make the model closer to real-life signalling situations, which is also the motivation in the present paper. Carlsson and Dasgupta (1997) eliminate unintuitive equilibria in the noiseless model by using noise and passing to the noiseless limit.

The experimental papers of Jeitschko and Normann (2012) and de Haan et al. (2011) use noisy one-shot signalling models with normally distributed noise. Their focus is on one-shot interactions and whether signalling is observed experimentally. The present paper studies repeated signalling theoretically.

Noisy effort exerted over time is also found in the career concerns literature that started from Holmström (1999). The difference of most of the career concerns models from the present paper is that the incomplete information is symmetric—the sender does not know his own type. The baseline signalling model also differs from Holmström (1999) in that effort is unproductive and the signal does not depend on type, but these assumptions are relaxed in the extensions.

The reputation models following Kreps and Wilson (1982) and Milgrom and Roberts (1982) are also related to the dynamic signalling literature. There are players choosing costly actions to influence the beliefs of other players. The change in beliefs leads to some benefit in the future. Most of the reputation literature features private values—players do not care about the types of other players, only their actions. This is the polar opposite of signalling. In signalling, it is the type of a player that is payoff-relevant to other players, not the action. Furthermore, in the present paper all types are strategic, unlike in most reputation models.

7 Conclusion

A tractable repeated noisy signalling model is proposed and the unique affine equilibrium is found. An advantage over the previous literature is the ease of describing nonstationary environments such as election campaigns, mating seasons, critical learning periods during development. Private instead of public monitoring of the signal by the market does not affect the results. Neither do some forms of competition between multiple signallers.

Closed form comparative statics are obtained, some of which are standard, some surprising. Among the latter, the most unexpected result is that when the cost difference between types changes, then generically there are types whose effort and payoff increase in their marginal cost.

Many extensions to the basic pure signalling model are possible, retaining tractability. Examples solved in this paper pertain to time-varying type, human capital accumulation, exogenous information revelation, productive effort and competition between signallers. Any subset of these can be combined without complications.

Future research will focus on incomplete information of the sender as well as the market. A seller may be uncertain about how much the buyers value quality. A monopolist may only imperfectly observe the demand. The anticipated consequence is that the sender experiments by choosing various (myopically suboptimal) signalling efforts over time and observing the resulting benefit.

References

Admati, A. R. and M. Perry (1987): "Strategic delay in bargaining," *The Review of Economic Studies*, 54, 345–364.

Alós-Ferrer, C. and J. Prat (2012): "Job market signaling and employer learning," *Journal of Economic Theory*, 147, 1787–1817.

Carlsson, H. and S. Dasgupta (1997): "Noise-Proof Equilibria in Two-Action Signaling Games," *Journal of Economic Theory*, 77, 432 – 460.

Cho, I.-K. and D. M. Kreps (1987): "Signaling games and stable equilibria," *The Quarterly Journal of Economics*, 179–221.

Daley, B. and B. Green (2012): "Waiting for News in the Market for Lemons," *Econometrica*, 80, 1433–1504.

——— (2014): "Market signaling with grades," *Journal of Economic Theory*, 151, 114–145.

de Haan, T., T. Offerman, and R. Sloof (2011): "Noisy signaling: theory and experiment," *Games and Economic Behavior*, 73, 402–428.

Dilme, F. (2014a): "Dynamic Noisy Signaling in Discrete Time," Working paper.

——— (2014b): "Dynamic Quality Signaling with Hidden Actions," Working paper 14-019, Penn Institute for Economic Research, Department of Economics, University of Pennsylvania.

Dilme, F. and F. Li (2014): "Dynamic Education Signaling with Dropout Risk," Working paper, Penn Institute for Economic Research, Department of Economics, University of Pennsylvania.

GRYGLEWICZ, S. (2009): "Signaling in a stochastic environment and dynamic limit pricing," mimeo, Tilburg University.

HEINSALU, S. (2014): "Noisy signalling over time," Working paper.

HERRNSTEIN, R. AND C. MURRAY (1994): *The Bell Curve: Intelligence and Class Structure in American Life*, Free Press.

HOLMSTRÖM, B. (1999): "Managerial incentive problems: A dynamic perspective," *The Review of Economic Studies*, 66, 169–182.

HÖRNER, J. AND N. VIEILLE (2009): "Public vs. private offers in the market for lemons," *Econometrica*, 77, 29–69.

JEITSCHKO, T. D. AND H.-T. NORMANN (2012): "Signaling in deterministic and stochastic settings," *Journal of Economic Behavior & Organization*, 82, 39–55.

KAYA, A. (2009): "Repeated signaling games," *Games and Economic Behavior*, 66, 841 – 854, special Section In Honor of David Gale.

KREMER, I. AND A. SKRZYPACZ (2007): "Dynamic signaling and market breakdown," *Journal of Economic Theory*, 133, 58–82.

KREPS, D. M. AND R. WILSON (1982): "Reputation and imperfect information," *Journal of Economic Theory*, 27, 253–279.

KYLE, A. S. (1985): "Continuous auctions and insider trading," *Econometrica: Journal of the Econometric Society*, 1315–1335.

LELAND, H. AND D. PYLE (1977): "Signaling and the valuation of unseasoned new issues," *Journal of Finance*, 37, 1–10.

MATTHEWS, S. A. AND L. J. MIRMAN (1983): "Equilibrium limit pricing: The effects of private information and stochastic demand," *Econometrica: Journal of the Econometric Society*, 981–996.

MCLENNAN, A., P. K. MONTEIRO, AND R. TOURKY (2014): "On the uniqueness of equilibrium in the Kyle model," CORE Discussion Papers 761, Escola de P'os-Graduaçã em Economia da Fundação Getulio Vargas.

MILGROM, P. AND J. ROBERTS (1982): "Predation, reputation, and entry deterrence," *Journal of Economic Theory*, 27, 280–312.

NÖLDEKE, G. AND E. VAN DAMME (1990): "Signalling in a dynamic labour market," *The Review of Economic Studies*, 57, 1–23.

RODDIE, C. (2012): "Signaling and Reputation in Repeated Games, II: Stackelberg limit properties," Working paper, University of Cambridge.

SPENCE, M. (1973): "Job Market Signaling," *The Quarterly Journal of Economics*, 87, pp. 355–374.

Swinkels, J. M. (1999): "Education signalling with preemptive offers," *The Review of Economic Studies*, 66, 949–970.

Vickers, J. (1986): "Signalling in a model of monetary policy with incomplete information," *Oxford Economic Papers*, 443–455.

Weiss, A. (1983): "A sorting-cum-learning model of education," *The Journal of Political Economy*, 420–442.