Noisy signalling over time

Sander Heinsalu *

October 4, 2014

Abstract

This paper examines signalling when the sender exerts effort and receives benefit over time. Receivers only observe a noisy public signal about effort, which has no intrinsic value.

Time introduces novel features to signalling. In some equilibria, a sender with a higher cost of effort exerts strictly more effort than his low-cost counterpart. Noise leads to robust predictions: pooling on no effort is always an equilibrium, while pooling on positive effort cannot occur. Whenever pooling is not the unique equilibrium, informative equilibria with a simple structure are shown to exist for some prior belief.

Keywords: Dynamic games, signalling, incomplete information, information economics.

JEL classification: D82, D83, C73.

1 Introduction

Most signalling situations involve noisy observation of the sender's effort. In many, effort is exerted over time. Such situations occur in education, finance, advertising, management, biology, anthropology and politics. For example, a politician may be a (relatively) honest or a corrupt type, and can signal honesty by following the law to the letter (paying taxes in full, refraining from speeding and bribe-taking). The cost of abiding by the law is incurred at all times. Voters learn of low effort only when some random events, such as scandals occur.

*School of Economics, University of Queensland.
Email: sander.heinsalu@yale.edu, website: http://sanderheinsalu.com/, address: Colin Clark Building Level 6, St Lucia, Brisbane, QLD 4072, Australia.
The author is greatly indebted to Johannes Hörner for many enlightening discussions about this research. Numerous conversations with Larry Samuelson have helped improve the paper. The author is grateful to Eduardo Faingold, Dirk Bergemann, Juuso Välimäki, Philipp Strack, Willemien Kets, Tadashi Sekiguchi, Flavio Toxvaerd, Françoise Forges, Jack Stecher, Daniel Barron, Benjamin Golub, Florian Ederer, Vijay Krishna, Sambuddha Ghosh, Sergiu Hart, Idione Meneghel, Jörgen Weibull, Alessandro Bonatti, Zvika Neeman, Scott Kominers, Andrzej Skrzypacz, George Mailath, Joel Sobel, three anonymous referees and participants of many conferences and seminars for comments and suggestions. Any remaining errors are the author's. Financial support from Yale University and the Cowles Foundation is gratefully acknowledged.

Results robust across noise structures are obtained. Pooling on zero effort is always an equilibrium, while pooling on positive effort cannot occur. There are always prior beliefs at which pooling is the unique equilibrium, and for some parameter values, pooling is unique for all prior beliefs. When pooling is not unique, simple informative equilibria exist for some prior belief. In addition to the robust results, novel dynamics appear. For example, there exist equilibria in which the high-cost type exerts strictly more effort initially than the low-cost type. Intuition from the previous literature may thus be misleading. The rich dynamics are due to the multiple opportunities for exerting effort, while the robust features are driven by the noise.

The environment is the natural adaptation of Spence (1973). The players are a sender and a competitive market of receivers. The sender is either a high-cost or a low-cost type. Type is private information. Receivers share a common prior belief about the type. The sender continuously chooses his effort level. Formally, the types only differ in their flow cost of effort, but low cost can be interpreted as high productivity, for which receivers pay more.

Receivers observe a noisy public signal about the effort, rather than the effort itself. The signal is modelled as a Poisson process. Effort either increases the intensity of the Poisson process (this is called the good news case), or decreases it (bad news). Unlike in the strategic experimentation literature, in this paper good news are a sign of high effort, which need not come from the low-cost type, and similarly for bad news. Whether a signal is good or bad news about the type is endogenous.

Examples of good news are awards and publications in education and research, attracting venture capital in finance, favourable coverage in a respected magazine or show in advertising, a famous businessperson joining the firm in management. Bad news are small antlers or dull plumage in biology, failing a test of strength or bravery in anthropology, a scandal in politics.

In the model, the sender derives a flow benefit directly from the posterior belief of the receivers. For example, a politician's chance of winning a surprise by-election, decision rights in the party and personal pride depend on the voters' current opinion of the politician.

Attention is restricted to Markov stationary equilibria. Pooling on no effort is an equilibrium for general noise structures. A sufficient condition is that whenever the receivers expect the sender's effort to be type-independent, beliefs stay constant regardless of signals. This condition is satisfied by Bayesian posterior beliefs. If beliefs do not respond to signals, then effort provides no benefit to the sender. This makes both types switch to zero effort to minimize cost. The preceding reasoning can also be used to rule out pooling on positive effort.

Pooling is the unique equilibrium when the benefit of signalling is too small to incentivize the low-cost type to signal. The cost of signalling is the same for all prior beliefs, while the benefit depends on the difference between the posterior beliefs after different signals. This difference is smaller for extreme beliefs, for a given imperfectly revealing signal structure. There are always prior beliefs high or low enough to make the benefit smaller than the cost.

Whenever pooling is not the unique equilibrium, there is a prior belief for

which informative equilibria with a simple structure exist. The high-cost sender never exerts effort. The low-cost sender initially exerts maximum effort, switching to zero effort when the belief becomes high or low enough. This can be interpreted as reputation building by the low-cost sender.

When the sender's benefit is concave in the receivers' belief, the high-cost sender strictly prefers pooling to any informative equilibrium. This is because in an informative equilibrium, the posterior belief has positive variance, and the high-cost type expects this belief to become less favourable on average over time.

In the bad news model, the high-cost sender exerts more effort than the low-cost in some equilibria. In these equilibria, effort reduces the likelihood of future information revelation. The high-cost sender's payoff decreases more due to information revelation, thus he has more incentive to avoid it by exerting effort. Intuitively, the corrupt do more to avoid transparency.

In the good news model, the high-cost sender exerts no more effort (at any belief) than the low-cost in all equilibria. The good news model most closely resembles Spence (1973). For example, the high-cost sender prefers pooling to all other equilibria. This is true even when the benefit is convex in the receivers' belief (so that learning has positive value). The reason is that when the equilibrium effort of the high-cost type is less than the maximum, then by the linearity of the cost and signal structure, the high-cost type is indifferent between equilibrium effort and no effort. No effort means no signals. By indifference, the equilibrium payoff of the high-cost type must be the same as in the absence of signals. In the absence of signals, belief drifts down under good news when the high-cost sender is expected to exert less effort than the low-cost. Flow benefit increases in belief, so a downward drift in belief lowers the payoff to the sender.

Continuous time signalling is also studied in Daley and Green (2012), Gryglewicz (2009) and Dilme (2014). In all three, the benefit of signalling is received in a lump sum when the sender decides to stop the game. In Daley and Green (2012), the signal process is exogenous. In Gryglewicz (2009), one type of the sender is a commitment type. In the present paper, the benefit is a flow, both types are strategic and effort controls the signal process.

The results for repeated noiseless signalling games of Kaya (2009) and Roddie (2012) resemble those of Spence (1973) and differ from the current paper in that pooling on positive effort is always an equilibrium and informative equilibria always exist.

Section 2 introduces the Poisson model and discusses equilibria in the bad news (Section 2.1) and the good news (Section 2.2) cases. Both bad and good news are added to the model in Section 2.3 and an exogenous lower bound on the signal rate is studied in Section 2.4. The related literature is discussed in more detail in Section 3.1. Additional results are provided in the online appendix, including models of signalling with Brownian noise.

2 Poisson signalling

This section presents the model where effort changes the intensity of a Poisson signal process. Both the good news and the bad news cases are considered, but first the setup of the model is described.

Time is continuous and the horizon is infinite. There is a strategic sender and a competitive market of receivers. The sender has two types, H and L, with initial log likelihood ratio[1] $l_0 \in \mathbb{R}$ that is common knowledge. The sender knows his type, the receivers do not. A generic log likelihood ratio l is an element of $\overline{\mathbb{R}} = \mathbb{R} \cup \{\infty, -\infty\}$. The log likelihood ratio corresponding to $\Pr(H) = 1$ is $l = \infty$ and corresponding to $\Pr(H) = 0$ is $l = -\infty$.

The sender has action set $[0, 1]$ at each instant of time. The action 0 is interpreted as no effort of signalling and the action 1 as maximal effort. Effort e costs type θ sender $c_\theta e$, with $c_L > c_H > 0$. Due to the linearity of cost and signal rate (discussed below) in effort, an alternative view is that the action set is $\{0, 1\}$, with mixing corresponding to $e \in (0, 1)$. To avoid the technical difficulties of defining behavioural strategies in continuous time, the action set is taken to be $[0, 1]$ and mixing is not allowed.

Effort benefits the sender via its effect on the signal process, which drives the market's log likelihood ratio process (l_t). The sender is assumed to derive flow benefit $\beta(l)$ directly from the market's log likelihood ratio l.[2] The function β is assumed strictly increasing, bounded and continuously differentiable. Denote the flow benefit from $l = \infty$ (corresponding to $\Pr(H) = 1$) by $\beta_{\max}$ and from $l = -\infty$ by $\beta_{\min}$. The function β is linear in the posterior belief if it has the form $\beta(l) = k_1 \frac{\exp(l)}{1+\exp(l)} + k_2$, with $k_1 > 0$, because the probability corresponding to log likelihood ratio l is $\frac{\exp(l)}{1+\exp(l)}$.

A *pure Markov stationary strategy* $(e_L, e_H) : \overline{\mathbb{R}} \to [0, 1]^2$ maps the log likelihood ratio to efforts of the types. To simplify the statements to follow, attention is restricted to e_L, e_H piecewise continuous[3] and at every discontinuity, continuous from the left or the right. The state variable l is the left limit l_{t-} of the log likelihood ratio (with the convention $l_{0-} = l_0$), so jumps are not anticipated by a strategy. Henceforth, only pure Markov stationary strategies are considered and the 'pure Markov stationary' phrase is omitted. The online appendix defines more general strategies and presents the signal process in greater detail.

Signals take the form of a Poisson process. In the good news (breakthrough) case, the Poisson rate of signals at time t is $e_t\lambda + \epsilon$, with $\epsilon \geq 0$. The parameter

[1] Throughout this paper, log likelihood ratio l is used instead of belief $\Pr(H) = \frac{\exp(l)}{1+\exp(l)}$, as this simplifies formulas in the dynamic models to follow. All results can be restated in terms of beliefs.

[2] This can be microfounded by assuming that each receiver has a unique one-shot best response $a^*(l)$ to each log likelihood ratio $l \in \mathbb{R}$. Since each receiver is infinitesimal, their current action does not influence the future, so in any equilibrium each receiver must play the one-shot best response. The sender is then assumed to derive flow benefit $\beta_a(a^*(l))$ from the receivers' action $a^*(l)$, with $\beta_a(a^*(\cdot)) = \beta(\cdot)$.

[3] A piecewise continuous strategy is understood to have at most finitely many discontinuities on $\overline{\mathbb{R}}$.

$\lambda \in (0, \infty)$ is interpreted as the informativeness of effort and e_t denotes the effort at t. Initially, the case $\epsilon = 0$ is studied. Section 2.4 looks at $\epsilon > 0$. In the bad news (breakdown) case, the rate of signals is $(1 - e_t)\lambda + \epsilon$, which decreases in effort. Note that with $\epsilon = 0$, zero effort in the good news case or maximal effort $e = 1$ in the bad news case ensures no signals occur. The main results continue to hold with $\epsilon > 0$, but the formulas become more complicated and in some cases, closed forms are no longer obtained.

The signals here depend on effort, not type, unlike in the strategic experimentation literature. Whether a signal of high effort is good or bad news about the type depends on the strategy expected from the sender.

The receivers observe the public signals, but not the sender's effort or type. Since the signal is public, the l it leads to is common to all receivers.

The Bayesian updating of the market's log likelihood ratio is described next. Denote the strategy the market expects by $e^* = (e_H^*, e_L^*)$. This notation is also used for equilibrium strategies. If a signal occurs at log likelihood ratio $l \in \mathbb{R}$, then in the good news case, l jumps to

$$j^g(l) = l + \ln\left(\frac{\lambda e_H^*(l) + \epsilon}{\lambda e_L^*(l) + \epsilon}\right).$$

In the bad news case, the log likelihood ratio jumps to

$$j^b(l) = l + \ln\left(\frac{\lambda(1 - e_H^*(l)) + \epsilon}{\lambda(1 - e_L^*(l)) + \epsilon}\right).$$

When $\epsilon = 0$, the fraction in the above formulas becomes $\frac{0}{0}$ for some e^*. In that case, the convention $\frac{0}{0} = 1$ is used, which comes from the limit of Bayesian updating as $\epsilon \to 0$. Also, when $\epsilon = 0$, the jumps $j^g(l), j^b(l)$ reach $-\infty$ or ∞ for some e^* for any l. Assume that once l reaches $\pm\infty$, it remains constant thereafter, so signals are ignored when there is certainty about the type (a perfectly informative signal occurred in the past). This is implied by Bayes' rule when $\epsilon > 0$, but when $\epsilon = 0$, Bayes' rule does not apply in some cases, e.g. if the signal has zero probability conditional on the (already certain) type. If Bayes' rule does not apply, there are many ways to specify beliefs, some of which change the results. This is discussed in Section 3 below. Belief updating based on Bayes' rule is summarized in the following lemma.

Lemma 1. *Under good news, the log likelihood ratio at t is*

$$l_t = l_0 - \lambda \int_0^t e_H^*(l_s) - e_L^*(l_s) ds + \sum_{k=1}^{n} \left[j^g(l_{\tau_k}) - l_{\tau_k}\right], \tag{1}$$

and under bad news, it is

$$l_t = l_0 + \lambda \int_0^t e_H^*(l_s) - e_L^*(l_s) ds + \sum_{k=1}^{n} \left[j^b(l_{\tau_k}) - l_{\tau_k}\right], \tag{2}$$

where (e_H^*, e_L^*) is the strategy the market expects and $\tau_1, \ldots, \tau_n$, with $0 \leq \tau_1 < \cdots < \tau_n \leq t$, are the times at which the signal occurs.

Given the signal times, the solution to (1) is the log likelihood ratio process $(l_t)_{t\geq 0}$ under good news, and the solution to (2) is the process under bad news.

Define the *reachable set* of log likelihood ratios

$$\mathcal{L}(e^*) = \left\{ l \in \overline{\mathbb{R}} : \exists s \in \mathbb{R}_+ \ \exists \text{ a path of } (l_t) \text{ s.t. } l_s = l \right\},$$

where l_t is given in Lemma 1. The reachable set depends on the strategy the market expects and on l_0. For each e^* and l_0, only behaviour on the reachable set is discussed subsequently. Since no deviation can take l outside the reachable set, behaviour outside it can be arbitrary.

Given the strategy $e^* = (e_L^*, e_H^*)$ the market expects, the payoff of type θ from the effort function $e_\theta(\cdot)$ and the log likelihood ratio process $(l_t)_{t\geq 0}$ is the expected discounted sum of flow payoffs

$$J_{l_0}^{e_\theta}(e^*) = \mathbb{E}^{e_\theta}\left[\int_0^\infty \exp(-rt)[\beta(l_t) - c_\theta e_\theta(l_t)]dt | l_{t=0} = l_0\right], \tag{3}$$

where the expectation is over the stochastic process $(l_t)_{t\geq 0}$, given $e_\theta(\cdot)$. The discount rate is $r > 0$.

Except for jumps, l evolves deterministically given the market expectations (e_L^*, e_H^*). Given l at the time of a jump, the size of the jump is deterministic.

Given a strategy e^* expected by the receivers, the payoff to type θ from the best response starting at l is denoted by $V_\theta(l)$. If the market expects a Markov stationary strategy, then every time l is reached, the continuation value of type θ from l is well defined and independent of the path of (l_t) that led to l. The dependence of $V_\theta(l)$ on e^* is suppressed in the notation.

Lemma 2. *$V_H(l) \geq V_L(l)$ $\forall l_0 \in \overline{\mathbb{R}}$ $\forall e^*$ $\forall l \in \mathcal{L}(e^*)$, with strict inequality if under the optimal (e_L, e_H) starting at l, there is a positive probability of reaching some $\hat{l}$ with $e_L(\hat{l}) > 0$. $\frac{\beta_{\min}}{r} \leq V_\theta(l) \leq \frac{\beta_{\max}}{r}$, with strict inequalities if $l \in \mathbb{R}$.*

Proof. $V_\theta(l)$ is bounded above by $\int_0^\infty \exp(-rt)\beta_{\max}dt = \frac{\beta_{\max}}{r} \in \mathbb{R}$ and below by $\frac{\beta_{\min}}{r} \in \mathbb{R}$.

$V_H(l)$ is greater than the payoff to H from imitating an optimal strategy of L starting from l. An optimal strategy gives L the value $V_L(l)$ from l. H can imitate an optimal strategy of L at a lower cost, getting the same benefit, so the imitation payoff to H is greater than $V_L(l)$.

If the set of histories where reaching some $\hat{l}$ satisfying $e_L(\hat{l}) > 0$ has positive probability starting from l under the optimal strategy, then H can imitate L at a strictly lower cost, getting the same benefit, so $V_H(l) > V_L(l)$. □

Definition 1. A Markov stationary equilibrium consists of a strategy $e^* = (e_H^*, e_L^*)$ of the sender and a log likelihood ratio process $(l_t)_{t\geq 0}$ s.t.

1. given $(l_t)_{t\geq 0}$, e_θ^* maximizes (3) over e_θ,
2. given e^*, $(l_t)_{t\geq 0}$ is derived from (1) under good news and (2) under bad news.

In the definition, the strategy is a function of the log likelihood ratio process, which depends on the strategy. On the one hand, this seeming circularity is unavoidable due to the fixed point nature of equilibrium. On the other hand, strategies are defined as functions on $\overline{\mathbb{R}}$ without reference to (l_t). To further reassure the reader of the innocuousness of the apparent circularity, the online appendix defines signal histories without reference to strategies or (l_t), then defines strategies as functions of signal histories and finally derives (l_t) from histories and the expected strategy.

Def. 1 implies that on the reachable set, behaviour is optimal from any point on. Therefore the equilibrium concept could also be called Markov perfect. Deviations to non-Markov strategies are feasible, but there is always a stationary best response when the receivers expect a stationary strategy. This is because every time a given l is reached, the future looks the same, so the same effort level is a best response.

Henceforth 'equilibrium' means a pure Markov stationary equilibrium. Call an equilibrium *extremal* when the equilibrium efforts only take values in $\{0, 1\}$. These are analogous to the pure equilibria of a static model, because the cost and the signal rate are linear in effort. Extremal efforts will be shown to imply that in an interval of log likelihood ratios, H exerts maximal effort, and outside the interval zero effort, while L never exerts effort. Exerting effort initially and then potentially stopping can be interpreted as reputation building by H. A *pooling* equilibrium is defined by $e_L^*(l) = e_H^*(l)\ \forall l \in \mathcal{L}(e^*)$. It exists by the following lemma.

Lemma 3. *(a)* $e_\theta(l) = 0$ *is the unique best response to* $e_H^*(l) = e_L^*(l)$ *for* $\theta = H, L$*,*

(b) the pooling equilibrium with $e_L^*(l_0) = e_H^*(l_0) = 0$ *exists for any* $l_0 \in \overline{\mathbb{R}}$*,*

(c) in any equilibrium for any $l \in \mathcal{L}(e^*)$*,* $e_L^*(l) = e_H^*(l) > 0$ *cannot occur.*

Proof. If $\forall l \in \mathcal{L}(e^*)\ \forall t \geq 0\ \forall s > 0$, the assumptions $e_H^*(l) = e_L^*(l)$ and $l_t = l$ imply $l_{t+s} = l_t$ (this is satisfied by (1) and (2)), then upon reaching any $\hat{l} \in \overline{\mathbb{R}}$ with $e_H^*(\hat{l}) = e_L^*(\hat{l})$, l remains at $\hat{l}$ forever regardless of effort, thus there is no benefit to signalling. Signalling is costly, so zero effort is the unique best response for both types.

If both types are expected to exert no effort at l_0, then zero effort is the unique best response. This proves the existence of the pooling equilibrium with $e_L^*(l_0) = e_H^*(l_0) = 0$.

If for some $l \in \mathcal{L}(e^*)$, the receivers expect $e_L^*(l) = e_H^*(l) > 0$, then both types choose no effort at l. □

The reasoning ruling out $e_L^*(l) = e_H^*(l) > 0$ is similar to Bagwell (1995). Even signals likely after a deviation to $e = 0$ and unlikely under the equilibrium strategy are ascribed to noise, so deviating is not punished. Due to this, both types deviate.

Lemma 3 shows that an extremal equilibrium always exists, because pooling on zero effort is an extremal equilibrium. The same reasoning and results hold

in a variety of noisy signalling games, some of which are discussed in the online appendix.

Pooling is the unique equilibrium if the benefit of signalling is low enough relative to the cost. It is proved below (Propositions 7 and 10) that if pooling is not the unique equilibrium, then there exists a prior log likelihood ratio for which a nonpooling extremal equilibrium exists.

Proposition 4. *If* $\frac{\beta_{\max}-\beta_{\min}}{r} \leq \frac{c_H}{\lambda}$*, then pooling is the unique equilibrium for any* $l_0 \in \overline{\mathbb{R}}$*.*

Proof. If $l = -\infty$ and a signal occurring at rate $e\lambda$ changes l to ∞, or if $l = \infty$ and a signal occurring at rate $(1-e)\lambda$ changes l to $-\infty$, then the marginal benefit of effort is $\lambda\left[\frac{\beta_{\max}}{r} - \frac{\beta_{\min}}{r}\right]$. The marginal cost to H is c_H. At $l = -\infty$ or $l = \infty$, l does not respond to signals, either by Bayes' rule or its limit. For $l \in \mathbb{R}$, $\frac{\beta_{\min}}{r} < V_\theta(l) < \frac{\beta_{\max}}{r}$, so the marginal benefit of effort at any $l \in \mathbb{R}$ is strictly lower than $\lambda\left[\frac{\beta_{\max}}{r} - \frac{\beta_{\min}}{r}\right]$. If $c_H \geq \lambda\left[\frac{\beta_{\max}}{r} - \frac{\beta_{\min}}{r}\right]$, then both types have the unique best response $e = 0$. □

A similar result to Proposition 4 holds when the signal structure is Brownian or the game is one-shot and noisy. These cases are described in the online appendix.

It is intuitive that when the sender's benefit is concave in the receivers' posterior belief, then L always prefers the pooling equilibrium to any other equilibrium. This is because for L in an informative equilibrium the expected posterior is lower than the prior, and with a concave benefit, the variance in the posterior is not beneficial.

Lemma 5. *If* $\hat{\beta}(\mu) = \beta\left(\ln\frac{\mu}{1-\mu}\right)$ *is concave in* μ*, then for any equilibrium* e^* *and for all* $l \in \mathcal{L}(e^*)$*,* $V_L(l) \leq \frac{\beta(l)}{r}$ *in the good or the bad news model.*

Proof. In the pooling equilibrium, $V_L(l) = \frac{\beta(l)}{r}$ $\forall l \in \mathcal{L}(e^*)$. In an informative equilibrium, $\mu = \Pr(H)$ drifts down in expectation for L. The flow benefit is increasing in the posterior: $\hat{\beta}'(\mu) > 0$. The posterior has positive variance. For a concave $\hat{\beta}$, variance does not raise the payoff of L. Therefore the continuation payoff from l in an informative equilibrium is lower than in pooling. □

This result also holds in all the signalling models discussed in the online appendix.

2.1 The bad news case

In the bad news model, there exist equilibria in which for some beliefs of the receivers the L type exerts higher effort than H, despite the uniformly higher marginal cost of effort. This distinguishes the bad news case from the previous literature on signalling. The result is reminiscent of the countersignalling of Feltovich et al. (2002), but the mechanism is quite different. In the bad news

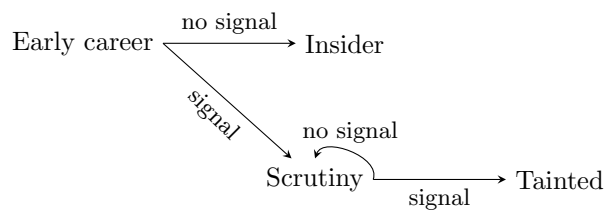

Figure 1: The four regimes in an equilibrium in which $e_L^* > e_H^*$.

model, it is the threat of future information revelation that incentivizes L to signal. This threat is not as severe for H.

The switched effort pattern $e_L^*(l) > e_H^*(l)$ is counterintuitive, because the benefit from a higher log likelihood ratio is the same for the types, but L has a higher marginal cost of signalling. A feature connected to this effort pattern is that V_H, V_L are nonmonotonic—they have a downward jump.

Example 1. Take $c_H = 0.1$, $c_L = 1.14$, $r = 1$, $\lambda = 2$, $\beta(l) = \frac{\exp(l)}{1+\exp(l)}$, $l_0 = 2 - \epsilon^2$, $\epsilon \in (0, 1)$. Then there exists an equilibrium in which in the interval $(2 - \epsilon, l_0]$, efforts are $e_L^* \in (0, 1)$ and $e_H^* = 0$. In the interval $[2, \infty)$, efforts are $e_L^* = 0$ and $e_H^* = 1$. Elsewhere, efforts are zero. Sufficient conditions for the existence of an equilibrium of this form are provided in the online appendix. The example satisfies these conditions, as can be calculated numerically.

The play in this equilibrium is illustrated in in Figure 1. To understand the equilibrium, think of a politician who can be a corrupt or an honest type and exerts effort to obey the law. Lawful behaviour reduces the frequency of scandals. The voters observe only the scandals and reward politicians believed to be more honest, voting for them with a higher probability if a surprise by-election should occur. The equilibrium displays four regimes, referred to as early career, insider, scrutiny and tainted. Play starts in the early career, during which the corrupt type (L) exerts positive effort and the honest type (H) no effort. If no scandal occurs by a given time, then the politician becomes an insider, which means that the voters ignore scandals and the politician no longer exerts effort. If instead a scandal occurs in the early career, then scrutiny results. Under scrutiny, H exerts maximum effort and L none. Under scrutiny, a scandal leads to a tainted reputation: voters are certain that the politician is corrupt and the politician exerts no effort.

The incentive for L to exert higher effort than H in the early career regime is created by the payoff difference between the insider and scrutiny regimes (the insider regime has a higher payoff). The payoffs to both types from the insider regime are equal, but the payoff from scrutiny is lower for L. The difference between insider and scrutiny payoffs is larger for L, so L can be incentivized to positive effort while H takes zero effort. Intuitively, future information revelation is worse for the L type, so L is willing to exert more effort in the early career to avoid it.

In the log likelihood ratio space, the four regimes of the switched effort

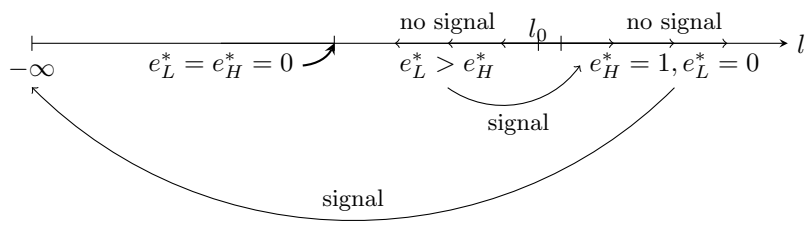

Figure 2: The log likelihood ratios in an equilibrium in which $e_L^* > e_H^*$.

equilibrium are depicted in Figure 2.

The low type exerting greater effort to avoid future information revelation has been noted in fiction by paragraphs and sentences such as "She made it to Seattle by five-thirty and to the museum by six, making sure to drive just under the speed limit. The last thing she needed was to get pulled over in a stolen truck with no license." (P. White "Undercover Stranger" p. 201) and "I drove just under the speed limit, not wanting to be stopped for speeding, maybe get frisked and be found with Tony's I.D." (P. S. Donoghue "Evil time" p. 190). The low type here is the criminal and the high type the innocent driver. Effort is driving slowly despite the desire to get to the destination fast, which is greater for a criminal needing to escape.

The effect of too high a signal lowering the belief of the receivers when the low type is expected to exert more effort is also well known in fiction: "Not even a parking ticket. No other citations. Too clean, as far as I'm concerned." (B. May "The Legacy of Lizzie Dolan" p. 130), "Did Riose refuse a bribe? Very suspicious; ulterior motives." (I. Asimov "Foundation and Empire" p. 83). The most famous related quote is "The lady doth protest too much, methinks." (W. Shakespeare "Hamlet" Act 3, Scene 2), often used to hint that vehement protestations of one's innocence are a sign of guilt.

In the educational environment, a switched effort equilibrium means that good students exert little effort early on in a course, knowing they can compensate for a failed test with future studying, while bad students start exerting effort initially and stop after a failure.

A switched effort equilibrium resembling Example 1 exists in the infinitely repeated noiseless signalling model of Kaya (2009), as discussed in the online appendix. That equilibrium is supported by belief threats off the equilibrium path. The switched effort equilibrium in the example above is robust to a positive lower bound on the signal rate, as shown at the end of Section 2.4. The positive lower bound implies all histories are on the equilibrium path.

The next section characterizes the set of extremal equilibria. There also exist equilibria in which $e_L^* \equiv 0$ and for an interval of reachable l, $e_H^*(l) \in (0, 1)$, otherwise $e_H^*(l) = 0$. These are described in the online appendix. These equilibria can be found numerically, but cannot be solved for in closed form.

2.1.1 Extremal equilibria

Extremal equilibria feature an interval of log likelihood ratios in which the high type exerts maximal effort and the low type zero. Outside the interval, both types exert no effort. This interval is called the *signalling region*. The prior log likelihood ratio is the lower boundary of the signalling region, because if $e_H^*(l) > e_L^*(l)$, then l drifts up.

Lemma 6. *If an equilibrium satisfies $e_\theta^*(l) \in \{0,1\}$ for $\theta = H, L$ for all $l \in \mathcal{L}(e^*)$, then $e_L^* \equiv 0$ and $\exists \bar{l} \geq l_0$ s.t. if $l \in [l_0, \bar{l})$, then $e_H^*(l) = 1$, and if $l \notin [l_0, \bar{l}]$, then $e_H^*(l) = 0$.*

Proof. Pooling on positive effort ($e_H^*(l) = e_L^*(l) = 1$) cannot occur. Maximal effort by L cannot occur, because with $e_L^*(l) = 1 > e_H^*(l)$, the jumps go to $l = \infty$, which is absorbing and gives the maximal payoff. However, L will not exert costly effort to avoid absorption in the state with the best possible payoff.

The log likelihood ratio cannot drift across a region on which $e_L^*(l) = e_H^*(l) = 0$. If $e_H^*(l) = 1$, $e_L^*(l) = 0$, then $j(l) = -\infty$, so there are no jumps into another region where $e_L^*(l) = 0$ and $e_H^*(l) = 1$. The set of log likelihood ratios on which the efforts of the types differ can only be an interval.

If $e_H^*(l) > e_L^*(l)$, then l drifts up, so l_0 is the lower boundary of any region where $e_H^*(l) > e_L^*(l)$. □

The signalling region cannot be closed at an infinite upper boundary $\bar{l} = \infty$, because at $l = \infty$, the log likelihood ratio does not respond to signals, so both types will choose $e = 0$. Closed signalling intervals of finite length can be replaced with intervals open on the right without changing the payoffs in any equilibrium. From now on, the signalling region is assumed open on the right and written as $[l_0, \bar{l})$.

The value functions are calculated next. At $\bar{l}$, the value functions of both types are $V_\theta(\bar{l}) = \frac{\beta(\bar{l})}{r}$. The boundary $\bar{l}$ can be infinite. In $[l_0, \bar{l})$, the value functions are solved for using Hamilton-Jacobi-Bellman (HJB) equations. A verification theorem (Theorem 4.6 in Presman et al. (1990) as modified for the discounted case in Yushkevich (1988)) is used to check that the solutions coincide with the value functions. The HJB equation of type θ is

$$rV_\theta(l) = \beta(l) + \lambda V_\theta'(l) + \max_e \left\{ \lambda(1-e)\left[\frac{\beta_{\min}}{r} - V_\theta(l)\right] - c_\theta e \right\}.$$

After reaching $\bar{l}$, incentives are trivial. In the signalling region, type θ chooses $e_\theta = 1$ if $-\lambda[\frac{\beta_{\min}}{r} - V_\theta(l)] - c_\theta \geq 0$. Rearranging this, one obtains the incentive constraints (ICs)

$$\frac{c_H}{\lambda} + \frac{\beta_{\min}}{r} \leq V_H(l), \qquad \frac{c_L}{\lambda} + \frac{\beta_{\min}}{r} \geq V_L(l), \tag{4}$$

which must hold for every l in the signalling region. These restrict the set of possible signalling regions and must be checked after solving for the candidate value functions.

To solve for the candidate value functions, substitute the equilibrium strategies $e_H^* = 1$ and $e_L^* = 0$ into the HJB equations of H and L. The HJB equations become the ordinary differential equations (ODEs)

$$\begin{aligned} rV_H(l) &= \beta(l) + \lambda V_H'(l) - c_H, \\ rV_L(l) &= \beta(l) + \lambda V_L'(l) + \frac{\lambda \beta_{\min}}{r} - \lambda V_L(l). \end{aligned}$$

In the absence of a signal, the log likelihood ratio rises continuously to $\bar{l}$. Assume $\bar{l}$ is finite (the case $\bar{l} = \infty$ is discussed in the online appendix). Then value matching gives $\lim_{l \to \bar{l}} V_\theta(l) = V_\theta(\bar{l}) = \frac{\beta(\bar{l})}{r}$, which provides the boundary condition for the ODEs. The solutions of the ODEs are

$$\begin{aligned} V_H(l) &= \exp\left(-r\frac{\bar{l}-l}{\lambda}\right)\frac{\beta(\bar{l})}{r} + \int_l^{\bar{l}} \frac{\beta(z) - c_H}{\lambda} \exp\left(-r\frac{z-l}{\lambda}\right) dz, \\ V_L(l) &= \exp\left(-(r+\lambda)\frac{\bar{l}-l}{\lambda}\right)\frac{\beta(\bar{l})}{r} \\ &\quad + \int_l^{\bar{l}} \left[\frac{\beta(z)}{\lambda} + \frac{\beta_{\min}}{r}\right] \exp\left(-(r+\lambda)\frac{z-l}{\lambda}\right) dz. \end{aligned} \tag{5}$$

These are continuously differentiable on $(l_0, \bar{l})$, with a right derivative at l_0 and a left derivative at $\bar{l}$, so by the verification theorem in Yushkevich (1988), they coincide with the candidate value functions. The ICs must be checked to confirm that the candidate value functions are indeed the value functions.

There exists an informative extremal equilibrium for some prior log likelihood ratio if the condition in Proposition 4 for pooling to be the unique equilibrium fails.

Proposition 7. *Suppose* $\frac{\beta_{\max} - \beta_{\min}}{r} > \frac{c_H}{\lambda}$. *Then* $\exists l_0 \in \mathbb{R}\ \exists \epsilon > 0$ *s.t. there exists a extremal equilibrium with signalling region* $[l_0, l_0 + \epsilon)$.

Proof. If $\frac{\beta_{\max} - \beta_{\min}}{r} > \frac{c_H}{\lambda}$, then $\exists l_0$ for which an informative extremal equilibrium can be constructed. Recall that β is strictly increasing. Define y by

$$y = \begin{cases} \infty & \text{if } \frac{\beta_{\max} - \beta_{\min}}{r} \leq \frac{c_L}{\lambda}, \\ \beta^{-1}\left(\frac{\beta_{\min}}{r} + \frac{c_L}{\lambda}\right) & \text{if } \frac{\beta_{\max} - \beta_{\min}}{r} > \frac{c_L}{\lambda}. \end{cases}$$

Take $\bar{l} \in \left(\beta^{-1}\left(\frac{\beta_{\min}}{r} + \frac{c_H}{\lambda}\right), y\right)$ in the extremal equilibrium, so that H has a strict incentive to signal at $\bar{l}$ and L has a strict incentive not to signal (recall that $V_H(\bar{l}) = V_L(\bar{l}) = \frac{\beta(\bar{l})}{r}$). By continuity and strict increasingness of V_H, V_L, $\exists \epsilon > 0$ s.t. $V_H(\bar{l} - \epsilon) \geq \frac{\beta_{\min}}{r} + \frac{c_H}{\lambda}$ and $V_L(\bar{l} - \epsilon) \leq \frac{\beta_{\min}}{r} + \frac{c_L}{\lambda}$, so H has an incentive to signal at $\bar{l} - \epsilon$ and L has an incentive not to signal. Take $l_0 = \bar{l} - \epsilon$. This completes the construction of the equilibrium. □

The intuition for Proposition 7 is that if the jumps from ∞ to $-\infty$ strictly incentivize H to signal, then there is an l_0 large enough s.t. jumps from l_0 to $-\infty$ do so as well. In that case there exists a extremal equilibrium with signalling region $[l_0, l_0 + \epsilon)$.

The final part of this section compares payoffs across extremal equilibria by varying $\bar{l}$. Based on (5), $V_L(l)$ and $V_H(l)$ are infinitely differentiable in $\bar{l}, r, \lambda, c_H, c_L$ for any $l \in \mathcal{L}(e^*)$, so derivatives can be used for comparative statics. Changing $\bar{l}$ changes e^* and therefore $\mathcal{L}(e^*)$. In that case, the comparison is between payoffs at an l that is in the reachable set both before and after changing $\bar{l}$.

Proposition 8. *If $\bar{l} \in \mathbb{R}$, then*

(a) $\frac{dV_L(l)}{d\bar{l}} > 0$ *iff* $\beta'(\bar{l}) - \beta(\bar{l}) + \beta_{\min} > 0$,

(b) $\frac{dV_H(l)}{d\bar{l}} > 0$ *iff* $\frac{\beta'(\bar{l})}{r} - \frac{c_H}{\lambda} > 0$.

Proof. The proof is by taking derivatives in (5). □

In Proposition 8, the effects of raising $\bar{l}$ on V_L are a higher payoff upon reaching $\bar{l}$ (the $\beta'(\bar{l})$ term), but a lower chance of reaching it (the $-\beta(\bar{l})$ term) and a higher chance of jumping to $l = -\infty$. The interpretation of the condition for $\frac{dV_H(l)}{d\bar{l}} > 0$ is that increasing $\bar{l}$ increases the payoff upon reaching $\bar{l}$ at a rate $\beta'(\bar{l})$ and increases the time during which the signalling cost is paid. Whether V_H increases or decreases in $\bar{l}$ depends on which effect dominates.

2.2 The good news case

The results about the good news model are presented next. In the good news case, the sender may be a firm trying to get endorsements from celebrities for its product, where an endorsement is less costly for a good quality firm and raises demand. Alternatively, a startup that manages to attract a well-known businessperson as a board member or manager sends a good signal about the quality of its business plan to potential investors, which expands financing opportunities. Or a researcher may exert effort to publish, with publications treated as signals of intelligence.

Under good news, the L type always prefers the pooling equilibrium, even when the flow benefit from the receivers' belief is convex. This is reminiscent of L preferring pooling to any informative equilibrium in noiseless models, but differs from the bad news case and from other noisy models considered in the online appendix.

Some preliminary observations about equilibrium efforts are collected in the following lemma. Since the lemma rules out $e_L^*(l) > e_H^*(l)$, switched effort equilibria do not exist in the good news case and the log likelihood ratio can only jump up: $j^g(l) \geq l$.

Lemma 9. *$\forall l_0 \in \overline{\mathbb{R}}\ \forall l \in \mathcal{L}(e^*)$, the equilibrium efforts satisfy $e_L^*(l) = e_H^*(l) = 0$ or $e_L^*(l) < e_H^*(l)$. Moreover, if $e_L^*(l) = 0 < e_H^*(l)$, then $e_H^*(l) = 1$.*

Proof. By Lemma 3, $e_L^*(l_t) = e_H^*(l_t) > 0$ cannot occur at any $l_t \in \mathcal{L}(e^*)$ in any equilibrium.

Three steps are needed to rule out $e_L^*(l) > e_H^*(l)$. First, in equilibrium L always has a best response that avoids jumps up. Second, in the region where $e_L^*(l) > e_H^*(l)$, $V_L(l)$ is bounded below by $\frac{\beta(l)}{r}$. Third, in the absence of jumps, V_L is increasing. If V_L is increasing, then L optimally does not exert effort to make l jump down.

Step 1. If $e_L^*(l) > e_H^*(l)$, then jumps go down. If $e_L^*(l) < e_H^*(l)$, then jumps go up, but $e_L^*(l) < 1$, so $e(l) = 0$ is a best response for L. No jumps occur with $e(l) = 0$.

Step 2. In the region where $e_L^*(l) > e_H^*(l)$, the log likelihood ratio drifts up and jumps down. Choosing $e(l) = 0$ avoids jumps and the flow cost. Starting at l, the flow benefit is at least $\beta(l)$ due to the upward drift.

Step 3. Consider $\hat{l} < l'$, with $e_L^*(l') > e_H^*(l')$. At $\hat{l}$, L has a best response that avoids jumps. $V_L(l') > \frac{\beta(l')}{r}$, because V_L is greater than the payoff to exerting zero effort forever. If the paths of l starting at $\hat{l}$ and l' never cross, then the flow benefit starting from $\hat{l}$ is always strictly below $\beta(l')$ and the cost is weakly higher. In that case, $V_L(\hat{l}) < V(l')$. If the paths of l starting at $\hat{l}$ and l' cross, then starting from l', the strategy that takes $e = 0$ until the paths cross and reverts to the optimal strategy thereafter yields a strictly higher payoff than the optimal strategy starting from $\hat{l}$. Before the crossing, the flow benefit starting from l' is strictly higher and the flow cost weakly lower than starting from $\hat{l}$. After the crossing, the payoffs are the same. As before, $V_L(\hat{l}) < V_L(l')$. This proves that $e_L^*(l) > e_H^*(l)$ cannot occur.

To rule out $e_L^*(l) = 0 < e_H^*(l) < 1$, suppose $\forall l \in (\underline{l}_1, l_0]$, $e_L^*(l) = 0 < e_H^*(l) < 1$ is expected. Since $e_H^*(l) \in (0,1)$, H is indifferent between $e = 0$ and $e = 1$, so switching H's choice from e_H^* to 0 on $(\underline{l}_1, l_0]$ does not change V_H. If $e(l) = 0\ \forall l \in (\underline{l}_1, l_0]$, then l drifts down deterministically to $\underline{l}_1$ and, if $\underline{l}_1 > -\infty$, reaches it and stops there forever. Consider $l', l'' \in (\underline{l}_1, l_0]$, with $l' > l''$. With $e(l) = 0\ \forall l \in (\underline{l}_1, l_0]$, starting at l' or l'', the flow cost is zero. Starting at l' initially yields a strictly higher flow benefit than starting at l'', and later (when $\underline{l}_1$ is reached) a weakly higher flow benefit. So V_H is strictly increasing in $(\underline{l}_1, l_0]$.

The jumps from $(\underline{l}_1, l_0]$ go to ∞, due to $e_H^*(l) > e_L^*(l) = 0$. If H is indifferent between $e(l^*) > 0$ and $e(l^*) = 0$ at some $l^* \in (\underline{l}_1, l_0]$, he is not indifferent at any $l \neq l^*$ in $(\underline{l}_1, l_0]$. This rules out $e_L^*(l) = 0 < e_H^*(l) < 1$ occurring over intervals of positive length in equilibrium. Efforts are continuous from the left or right in l, so the situation where $e_L^*(l) = 0 < e_H^*(l) < 1$ at only one point cannot occur. □

Restricting attention to extremal equilibria,[4] Lemma 9 implies that the only possible effort combinations are $e_H^*(l) = e_L^*(l) = 0$ and $e_H^*(l) = 1$, $e_L^*(l) = 0$. By a similar reasoning to the bad news case, an extremal equilibrium consists of an interval (called the signalling region) on which $e_H^*(l) = 1$, $e_L^*(l) = 0$ and

[4] Focussing on extremal equilibria is a restriction. There exist equilibria where $e_L^* \in (0,1)$ and $e_H^* = 1$, as shown in the online appendix.

outside which $e_H^*(l) = e_L^*(l) = 0$. In the good news case, the prior log likelihood ratio is the upper boundary of the signalling region. At the other boundary, the signalling region is assumed open, as in the bad news case. Denote the signalling region $(\underline{l}, l_0]$.

As in the bad news case, the value functions are calculated first. Then bounds on the set of signalling regions are discussed, followed by the existence of informative extremal equilibria. The reasoning and results are similar to the bad news case. Comparative statics are derived at the end of the section. These differ from the bad news model.

Outside $(\underline{l}, l_0]$, the value functions of both types are $V_\theta(l) = \frac{\beta(l)}{r}$. In $(\underline{l}, l_0]$, the value functions are solved for using the HJB equation and a verification theorem. The HJB equation for type θ is

$$rV_\theta(l) = \beta(l) - \lambda V_\theta'(l) + \max_e e\left\{\lambda\left[\frac{\beta_{\max}}{r} - V_\theta(l)\right] - c_\theta\right\}.$$

If $e_L^*(l) = e_H^*(l) = 0$, then incentives are trivial. At every l in the signalling region, the incentive constraints

$$\lambda\left[\frac{\beta_{\max}}{r} - V_H(l)\right] - c_H \geq 0, \qquad \lambda\left[\frac{\beta_{\max}}{r} - V_L(l)\right] - c_L \leq 0$$

must be satisfied in order for H to choose $e_H(l) = 1$ and L to choose $e_L(l) = 0$. These incentive constraints restrict the set of possible signalling regions and must be checked after solving for the candidate value functions.

The constraints have a simple interpretation: the marginal benefit of an increase in effort is the increased probability of jumping to $l = \infty$ and getting $\beta_{\max}$ forever instead of the current value $V_\theta(l)$. This probability increases in effort at rate λ. The marginal cost of effort is c_θ. If marginal cost minus marginal benefit is positive, type θ chooses $e = 1$, otherwise $e = 0$.

To solve for the candidate value functions, substitute the equilibrium strategies $e_H^*(l) = 1$ and $e_L^*(l) = 0$ into the HJB equations of H and L. The HJB equations become the ODEs

$$\begin{aligned}
&\lambda V_H'(l) + (\lambda + r)V_H(l) = \beta(l) + \frac{\lambda\beta_{\max}}{r} - c_H, \\
&\lambda V_L'(l) + rV_L(l) = \beta(l).
\end{aligned}$$

In the absence of a signal, the log likelihood ratio falls continuously to $\underline{l}$. For $\underline{l} > -\infty$, the value matching condition $\lim_{l\to\underline{l}} V_\theta(l) = V_\theta(\underline{l}) = \frac{\beta(\underline{l})}{r}$ holds, because close to $\underline{l}$, reaching it is likely and a jump to another value unlikely. The limit gives the boundary condition $V_\theta(\underline{l}) = \frac{\beta(\underline{l})}{r}$ for the ODEs. The case where $\underline{l} = -\infty$ is discussed in the online appendix.

The solutions of the ODEs are

$$
\begin{aligned}
V_H(l) &= \exp\left(-(r+\lambda)\frac{l-\underline{l}}{\lambda}\right)\frac{\beta(\underline{l})}{r} \\
&\quad + \int_{\underline{l}}^{l}\left[\frac{\beta(z)-c_H}{\lambda}+\frac{\beta_{\max}}{r}\right]\exp\left(-(r+\lambda)\frac{l-z}{\lambda}\right)dz, \qquad (6)\\
V_L(l) &= \exp\left(-r\frac{l-\underline{l}}{\lambda}\right)\frac{\beta(\underline{l})}{r}+\int_{\underline{l}}^{l}\frac{\beta(z)}{\lambda}\exp\left(-r\frac{l-z}{\lambda}\right)dz.
\end{aligned}
$$

The solutions to the HJB equations of the types are continuously differentiable on $(\underline{l}, l_0)$, with a right derivative at $\underline{l}$ and a left derivative at l_0, so by the verification theorem in Yushkevich (1988), the solutions are the candidate value functions. If the ICs are satisfied in $(\underline{l}, l_0]$, then V_H, V_L are the value functions.

The parameter values for which there exists a prior log likelihood ratio permitting an informative extremal equilibrium to exist are given in Proposition 10.

Proposition 10. *Suppose* $\frac{\beta_{\max}-\beta_{\min}}{r} > \frac{c_H}{\lambda}$. *Then* $\exists l_0 \in \mathbb{R}\ \exists \epsilon > 0$ *s.t. there exists an extremal equilibrium with signalling region* $[l_0, l_0+\epsilon)$.

Proof. If $\frac{\beta_{\max}-\beta_{\min}}{r} > \frac{c_H}{\lambda}$, then for some l_0 an extremal equilibrium can be constructed. Define y as follows. If $\frac{\beta_{\max}-\beta_{\min}}{r} \leq \frac{c_L}{\lambda}$, then set $y = -\infty$, otherwise set $y = \beta^{-1}\left(\frac{\beta_{\max}}{r} - \frac{c_L}{\lambda}\right)$.

Take $\underline{l} \in \left(y, \beta^{-1}\left(\frac{\beta_{\max}}{r} - \frac{c_H}{\lambda}\right)\right)$ in the extremal equilibrium, so that if $j(\underline{l}) = \infty$, then H has a strict incentive to signal at $\underline{l}$ and L has a strict incentive not to signal (recall that $V_H(\underline{l}) = V_L(\underline{l}) = \frac{\beta(\underline{l})}{r}$). By continuity and strict increasingness of V_H, V_L, $\exists \epsilon > 0$ s.t. $V_H(\underline{l}+\epsilon) \leq \frac{\beta_{\max}}{r} - \frac{c_H}{\lambda}$ and $V_L(\underline{l}+\epsilon) \geq \frac{\beta_{\max}}{r} - \frac{c_L}{\lambda}$, so H has an incentive to signal at $\underline{l}+\epsilon$ and L has an incentive not to signal. Take $l_0 = \underline{l}+\epsilon$. This completes the construction of an extremal equilibrium. □

Payoffs in extremal equilibria are compared next. The method is the same as in the bad news model, but the results differ. Based on (6), $V_H(l)$ and $V_L(l)$ are infinitely differentiable in $\underline{l}, r, \lambda, c_H, c_L$ for all $l \in \mathcal{L}(e^*)$, so derivatives can be used. Changing $\underline{l}$ changes e^* and therefore $\mathcal{L}(e^*)$. In that case, the comparison is between payoffs at an l that is in the reachable set both before and after changing $\underline{l}$.

Proposition 11. *If* $\underline{l} \in \mathbb{R}$, *then*

(a) $\frac{dV_L(l)}{d\underline{l}} > 0$,

(b) $\frac{dV_H(l)}{d\underline{l}} > 0$ *iff* $\beta'(\underline{l}) - \beta_{\max} + \beta(\underline{l}) + \frac{c_H r}{\lambda} > 0$,

Proof. The proof is by taking derivatives in (6). □

Comparing extremal equilibria, $\frac{dV_L(l)}{d\underline{l}} > 0$, so pooling gives L the highest payoff even when the benefit from the receivers' log likelihood ratio is arbitrarily

convex. The reason is that $e_L^* = 0$ in extremal equilibria, so L never receives good signals. In informative equilibria, there is a downward drift in l, which lowers V_L below pooling. This result contrasts with the bad news case where for β convex enough, L prefers information revelation to pooling.

The result that pooling always gives L the highest payoff in the good news model holds not just for extremal equilibria. By Lemma 9, a non-extremal equilibrium must feature $e_L^* \in (0, 1)$ in the signalling region, i.e. L is indifferent to receiving good signals and paying the signalling cost. In that case, $e = 0$ is still a best response for L, so V_L is unchanged if the chosen action of L is switched to 0, keeping expectations e_L^*, e_H^* equal to the equilibrium strategies. In other words, the payoff of L in an informative equilibrium is the same as it would be without good signals and with zero signalling cost. Due to the downward drift of l, this is lower than the pooling payoff.

The condition for $\frac{dV_H(l)}{d\underline{l}} > 0$ has a straightforward interpretation. Increasing $\underline{l}$ increases the payoff upon reaching $\underline{l}$ at a rate $\beta'(\underline{l})$, lowers the chance of jumping to $l = \infty$ (the $-\beta_{\max}$ term), increases the chance of reaching $\underline{l}$ (the $\beta(\underline{l})$ term) and reduces the time during which the signalling cost is paid. The balance of these effects determines whether V_H increases or decreases in $\underline{l}$. The result that H may prefer more or less information revelation depending on the parameters is the same as in the bad news model, but the reasons differ, because the effects underlying the $-\beta_{\max}$ and $\beta(\underline{l})$ terms in Proposition 11 are absent in the bad news model.

The next two sections relax some of the assumptions in the Poisson signalling model. Section 2.3 allows for good and bad news together and Section 2.4 studies the case where no effort level can make the signal rate zero.

2.3 Both good and bad news

In some circumstances, both good and bad news signals may occur. A firm's marketing expenditures increase the positive and decrease the negative media coverage, both of which influence demand. Similarly, lobbying targets both favourable and unfavourable news, which have an effect on policy. In education, good and bad grades both depend on the amount of studying.

If the Poisson rate of good news signals occurring is $\lambda_g e$, the rate of bad news is $\lambda_b(1-e)$ and the efforts the market expects from the types are e_L^* and e_H^*, then the HJB equation for type θ is

$$\begin{aligned} rV_\theta(l) &= \beta(l) + (\lambda_b - \lambda_g)(e_H^*(l) - e_L^*(l))V_\theta'(l) \\ &+ \max_e \{\lambda_g e \left[V_\theta \left(l + \ln \frac{e_H^*(l)}{e_L^*(l)} \right) - V_\theta(l) \right] \\ &+ \lambda_b(1-e) \left[V_\theta \left(l + \ln \frac{1-e_H^*(l)}{1-e_L^*(l)} \right) - V_\theta(l) \right] - c_\theta e\}. \end{aligned}$$

Type θ chooses $e_\theta = 1$ if

$$\lambda_g V_\theta\left(l + \ln\frac{e_H^*(l)}{e_L^*(l)}\right) - \lambda_g V_\theta(l)$$
$$- \lambda_b V_\theta\left(l + \ln\frac{1-e_H^*(l)}{1-e_L^*(l)}\right) + \lambda_b V_\theta(l) - c_\theta > 0$$

and only if a weak inequality holds.

The focus is on extremal equilibria in which type L always chooses $e = 0$ and type H chooses 1 in an interval $[l_0, \bar{l})$ or $(\underline{l}, l_0]$ of log likelihood ratios and 0 elsewhere. For simplicity, assume $\beta_{\max} > \frac{\lambda_b\beta_{\min} - \lambda_g\beta_{\max}}{\lambda_b - \lambda_g} + \frac{c_L r}{\lambda_b - \lambda_g}$ and $\beta_{\min} < \frac{\lambda_g\beta_{\max} - \lambda_b\beta_{\min}}{\lambda_g - \lambda_b} - \frac{c_L r}{\lambda_g - \lambda_b}$, which ensures $\underline{l} > -\infty$ and $\bar{l} < \infty$. In the signalling region, the jump after a good signal is to $l = \infty$ and the jump after a bad signal to $l = -\infty$. When $e_H^* = 1$, $e_L^* = 0$ is expected by the receivers, then the HJB equation in the signalling region is

$$rV_\theta(l) = \beta(l) + (\lambda_b - \lambda_g)V_\theta'(l) + \lambda_b\frac{\beta_{\min}}{r} - \lambda_b V_\theta(l)$$
$$+ \max_e e\left\{\lambda_g\frac{\beta_{\max}}{r} - \lambda_b\frac{\beta_{\min}}{r} + (\lambda_b - \lambda_g)V_\theta(l) - c_\theta\right\}.$$

Type θ chooses $e_\theta = 1$ if $\lambda_g\frac{\beta_{\max}}{r} - \lambda_b\frac{\beta_{\min}}{r} + (\lambda_b - \lambda_g)V_\theta(l) - c_\theta > 0$ and only if a weak inequality holds.

If $\lambda_b > \lambda_g$, then in the absence of a signal, l drifts up and eventually reaches $\bar{l}$. The signalling region has the form $[l_0, \bar{l})$. This implies $\lim_{l\to\bar{l}} V_\theta(l) = \frac{\beta(\bar{l})}{r}$, as in the bad news case.

If $\lambda_b < \lambda_g$, then in the absence of a signal, l drifts down and $\lim_{l\to\underline{l}} V_\theta(l) = \frac{\beta(\underline{l})}{r}$, as in the good news case. The signalling region has the form $(\underline{l}, l_0]$.

Substituting $e_H = 1$ and $e_L = 0$ into the HJB equations of the types, these become the ODEs

$$rV_H = \beta(l) + (\lambda_b - \lambda_g)V_H'(l) + \lambda_g\frac{\beta_{\max}}{r} - \lambda_g V_H(l) - c_H,$$
$$rV_L = \beta(l) + (\lambda_b - \lambda_g)V_L'(l) + \lambda_b\frac{\beta_{\min}}{r} - \lambda_b V_L(l).$$

The boundary condition depends on whether $\lambda_g > \lambda_b$ or vice versa. If $\lambda_b > \lambda_g$, then the boundary condition is $V_\theta(\bar{l}) = \frac{\beta(\bar{l})}{r}$ and the solutions are

$$V_H(l) = \exp\left(-(r+\lambda_g)\frac{\bar{l}-l}{\lambda_b - \lambda_g}\right)\frac{\beta(\bar{l})}{r}$$
$$+ \int_l^{\bar{l}}\left[\frac{\beta(z) - c_H}{\lambda_b - \lambda_g} + \frac{\lambda_g\beta_{\max}}{r(\lambda_b - \lambda_g)}\right]\exp\left((-(r+\lambda_g)\frac{z-l}{\lambda_b - \lambda_g}\right)dz,$$

$$V_L(l) = \exp\left(-(r+\lambda_b)\frac{\bar{l}-l}{\lambda_b-\lambda_g}\right)\frac{\beta(\bar{l})}{r}$$
$$+\int_l^{\bar{l}}\left[\frac{\beta(z)}{\lambda_b-\lambda_g}+\frac{\lambda_b\beta_{\min}}{r(\lambda_b-\lambda_g)}\right]\exp\left(-(r+\lambda_b)\frac{z-l}{\lambda_b-\lambda_g}\right)dz.$$

If $\lambda_b < \lambda_g$, the boundary condition is $V_\theta(\underline{l}) = \frac{\beta(\underline{l})}{r}$ and the solutions are

$$V_H(l) = \exp\left(-(r+\lambda_g)\frac{l-\underline{l}}{\lambda_g-\lambda_b}\right)\frac{\beta(\underline{l})}{r}$$
$$+\int_{\underline{l}}^{l}\left[\frac{\beta(z)-c_H}{\lambda_g-\lambda_b}+\frac{\lambda_g\beta_{\max}}{r(\lambda_g-\lambda_b)}\right]\exp\left(-(r+\lambda_g)\frac{l-z}{\lambda_g-\lambda_b}\right)dz,$$

$$V_L(l) = \exp\left(-(r+\lambda_b)\frac{l-\underline{l}}{\lambda_g-\lambda_b}\right)\frac{\beta(\underline{l})}{r}$$
$$+\int_{\underline{l}}^{l}\left[\frac{\beta(z)}{\lambda_g-\lambda_b}+\frac{\lambda_b\beta_{\min}}{r(\lambda_g-\lambda_b)}\right]\exp\left(-(r+\lambda_b)\frac{l-z}{\lambda_g-\lambda_b}\right)dz.$$

The set of extremal equilibria with finite signalling regions is similar to the case of only good news when $\lambda_g > \lambda_b$ and to the case of only bad news when $\lambda_g < \lambda_b$. This is because the candidate value functions and incentive constraints are similar.

In the knife-edge case of $\lambda_g = \lambda_b$, the log likelihood ratio stays constant at l_0 in the absence of signals, so $V_H(l_0) = \frac{\beta(l_0)-c_H}{r+\lambda_g} + \frac{\lambda_g\beta_{\max}}{r(r+\lambda_g)}$ and $V_L(l_0) = \frac{\beta(l_0)}{r+\lambda_b} + \frac{\lambda_b\beta_{\min}}{r(r+\lambda_b)}$.

The result that L always prefers pooling in the good news model is not sensitive to the rate of bad news being positive, provided it is less than the rate of good news. The intuition of why pooling gives L the highest payoff is that the rate of good signals is zero for L, because L exerts zero effort. From the viewpoint of L, in an informative extremal equilibrium, the log likelihood ratio can only decrease. This decreases the payoff to L, regardless of the curvature of β.

A low rate of good news compared to bad news ($\lambda_g \ll \lambda_b$) still permits switched effort equilibria to exist. Intuitively, all the incentive constraints in switched effort equilibria are strict except L's in the region where L mixes. If the game is perturbed a little, the strict incentive constraints still hold. The mixing probability expected from L can be changed to restore the indifference of L (the mixing probability expected determines the size of the jump in l, thus the value after the jump). Being again indifferent, L can mix with the probability expected.

2.4 Signal rate uniformly bounded below

This section looks at the case where even with zero effort in the good news case and even with maximal effort in the bad news case, the Poisson rate of signals is positive.

There are no zero-probability histories, so no refinements are needed. No signal is perfectly revealing, so even if efforts $e_L^* = 0$, $e_H^* = 1$ are expected, the jumps no longer go to ∞ or $-\infty$.

Call a log likelihood ratio process *boundedly informative* if $\forall e^*$ $\forall e$ $\forall \varepsilon > 0$ $\forall T > 0$ $\exists K > 0$ s.t. if $|l_t| > K$, then

$$\Pr\left(\left|\frac{\exp(l_{t+T})}{1+\exp(l_{t+T})} - \frac{\exp(l_t)}{1+\exp(l_t)}\right| < \varepsilon\right) > 1 - \varepsilon.$$

Intuitively, bounded informativeness ensures that close enough to certainty, l is expected to move arbitrarily slowly. If the Poisson rate of signals is uniformly bounded below by $\epsilon > 0$, then the log likelihood ratio process is boundedly informative.

For any boundedly informative l process, for $|l|$ large enough (equivalently, belief close enough to 0 or 1), pooling is the unique equilibrium. This is proved by two observations. First, due to discounting, payoffs more than T amount of time in the future have a small impact on the present. Second, due to $\beta \in C^2$ and bounded, if l_{t+T} stays "close" to l_t, then the movement of l has a small impact on the benefit. A large $|l|$ therefore makes the expected benefit of signalling small for any strategy expected by the receivers. With a sufficiently small expected benefit, both types optimally choose no effort.

In the good news case, the HJB equation of type θ under the expectations of the market e_L^*, e_H^* is

$$\begin{aligned} rV_\theta(l) &= \beta(l) - \lambda(e_H^*(l) - e_L^*(l))V_\theta'(l) \\ &+ \max_e \left\{(\lambda e + \epsilon)\left[V_\theta\left(l + \ln\frac{\lambda e_H^*(l)+\epsilon}{\lambda e_L^*(l)+\epsilon}\right) - V_\theta(l)\right] - c_\theta e\right\}. \end{aligned}$$

Interval equilibria are discussed next. Any signalling interval must be finite ($\underline{l} > -\infty$ and $\bar{l} < \infty$), because the signal structure is boundedly informative. Assume ϵ is small enough for the jumps to end in the pooling region (below $\underline{l}$ or above $\bar{l}$). Then the HJB equation becomes

$$rV_\theta(l) = \beta(l) - \lambda V_\theta'(l) + \max_e \left\{(\lambda e + \epsilon)\left[\frac{\beta(j(l))}{r} - V_\theta(l)\right] - c_\theta e\right\},$$

which can be rearranged as

$$\begin{aligned} (r+\epsilon)V_\theta(l) &= \beta(l) + \frac{\epsilon\beta(j(l))}{r} \\ &- \lambda V_\theta'(l) + \max_e \left\{\lambda e\left[\frac{\beta(j(l))}{r} - V_\theta(l)\right] - c_\theta e\right\}. \end{aligned}$$

This is the same as the HJB equation for an interval equilibrium of a game with $\epsilon = 0$, discount rate $r + \epsilon$ and flow benefit function $\beta(l) + \frac{\epsilon\beta(j(l))}{r}$. A similar rearrangement can be carried out for the bad news case. Thus a small positive rate of signals even at extreme effort levels does not greatly change the value

functions in interval equilibria. The comparative statics of equilibrium payoffs remain similar to the model where the signals can be perfectly informative.

What change with boundedly informative signals are the ICs—it is no longer the case that if $e_L^*(l) = 0, e_H^*(l) = 1$, then $V_\theta(j(l))$ is the same after every jump and independent of the type. Due to this, the ICs may bind at any l in the signalling region, not just at the endpoints. As a consequence, the set of bounds of interval equilibria need not be an interval any more.

The payoff comparison of pooling and informative equilibria for L in the good news case changes. It is no longer the case that L always prefers pooling, because of the $\frac{\epsilon\beta(j(l))}{r}$ term in the flow payoff. If $e_H^*(l) > e_L^*(l)$ is expected, then $j(l) > l$ and $\frac{\beta(j(l))}{r} > \frac{\beta(l)}{r}$. For any $\epsilon > 0$, there is a β with $\beta(\ln \frac{\mu}{1-\mu})$ convex enough in μ such that the payoff increase from the jumps in an informative equilibrium outweighs the payoff decrease from the downward drift of l. This makes L's payoff in an informative equilibrium higher than the pooling payoff $\frac{(1+\epsilon)\beta(l)}{r}$. However, for any β, there is a small enough $\epsilon > 0$ such that L prefers pooling to any informative equilibrium.

In the good news case, if ϵ is small enough, then $e_L^*(l) > e_H^*(l)$ still cannot occur for any reachable l. The proof idea is the same as with $\epsilon = 0$. The crucial comparison in that proof is strict and continues to hold with $\epsilon > 0$ small enough, because V_θ differs little between the case of $\epsilon > 0$ small enough and the case of $\epsilon = 0$.

The equilibria in the bad news case featuring $e_L^*(l) > e_H^*(l)$ for a nonempty interval of l still exist with $\epsilon > 0$. A numerical example that satisfies the sufficient conditions has parameters $c_H = 0.1$, $c_L = 1.14$, $r = 1$, $\lambda = 2$, $\epsilon = 0.001$ and $l_0 = 2.499999$. The benefit function is $\beta(l) = \frac{\exp(l)}{1+\exp(l)}$. The scrutiny region is $[2.5, 8)$ and the switched effort region is $(2.499, l_0]$. As in the preceding section on both good and bad news, the intuition for the survival of switched effort equilibria under small perturbations of the game is that all the incentives are either strict or can be restored by changing the mixing probability expected from L.

3 Discussion

The qualitative results that the good and bad news models have in common are that pooling on zero effort always exists, pooling on positive effort cannot occur, an extremal effort equilibrium exists for some prior whenever pooling is not unique, and the low type prefers pooling to any other equilibrium whenever the benefit from the receivers' belief is concave. These features are present in other noisy signalling models (discussed in the online appendix) as well, regardless of whether the game is one-shot or dynamic. The reasoning underlying these results is quite general and robust.

The switched effort equilibria in the bad news model depend on the particular signal structure. They do not exist in the good news case or the Brownian signalling model in the online appendix. The failure of single crossing in type

and effort despite the single crossing in type and cost is easier to generate in the noiseless case, as shown in the discussion of the repeated signalling model of Kaya (2009) in the online appendix. Many belief threats after histories off the equilibrium path become available in the absence of noise and can incentivize a variety of behaviour. Noise can force Bayes' rule to apply after all histories, as in Section 2.4, which reduces the set of equilibria but does not eliminate switched efforts.

If some effort levels can make the signal rate zero, as initially assumed in this paper, then there are histories off the equilibrium path, following which it is possible to impose different refinements on beliefs, e.g. the deviator is believed to be the type who gains the most from deviating. Some of these alternative belief specifications change the results. If the refinement in the good news model is that a zero-probability good signal reveals the high type, then pooling on zero effort is not an equilibrium for all parameter values. There is an incentive to exert effort if both types are expected to exert no effort. However, this incentive is not continuous in the expected efforts. If the high type is expected to exert arbitrarily small positive effort and the low type zero, then the incentive to signal is arbitrariliy small, while the incentive when both types are expected to exert zero effort is bounded away from zero under the refinement.

In the bad news model, the refinement that a zero-probability bad signal reveals the low type makes pooling on maximal effort an equilibrium for some parameter values. The reasoning is similar to the good news case in the previous paragraph, and a similar discontinuity occurs. Refinement issues only arise if the lower bound on the signal rate is zero. The positive lower bound in Section 2.4 leaves no scope for refinements and serves as a robustness check.

3.1 Literature

There are many papers that bear similarities to this one. This is not surprising, because the fact that in many cases signalling takes time was pointed out in Weiss (1983) and Admati and Perry (1987) already. Noisy signalling, on the other hand, was studied in Matthews and Mirman (1983).

Discrete-time repeated signalling is studied in Kaya (2009) and Roddie (2012). They differ from the present paper by the absence of noise in the observation of the sender's action. Kaya (2009) focusses on least-cost separating equilibria. Roddie (2012) provides general conditions for reputation effects to arise. Neither in Kaya (2009) nor in Roddie (2012) do the existence and payoffs of separating equilibria depend on the prior, unlike in the present paper. A more detailed overview of Kaya (2009) is given in the online appendix.

Signalling over time has also been studied in continuous time. Daley and Green (2012), Gryglewicz (2009) and Dilme (2014) use Brownian noise to model imperfect observation of the sender's action. In Dilme (2014), the sender (an entrepreneur) decides how much costly effort to exert over time, as well as when to stop the game (sell the firm) and receive a final benefit. This contrasts with the present paper, in which benefit accumulates continuously. In Daley and Green (2012), the uninformed traders receive information (observations of a

diffusion process) exogenously over time and the informed trader decides when to stop the game (execute the trade) and receive a final payoff. Gryglewicz (2009) looks at limit pricing over time. The low-cost incumbent is a commitment type and the high-cost incumbent decides when to stop imitating the low-cost type. Unlike in Daley and Green (2012), the signal is endogenous in the present paper, and unlike Gryglewicz (2009), both types are strategic.

Less closely related works on repeated noiseless signalling are Nöldeke and van Damme (1990) and Swinkels (1999). In these, the sender pays the signalling cost first, and receives the benefit only upon deciding to stop signalling forever. In the current paper, the benefit is received concurrently with the payment of the cost. Nöldeke and van Damme (1990) find a unique informative equilibrium. Using different informational assumptions, Swinkels (1999) finds a unique pooling equilibrium. The models in the current paper have many informative equilibria and one pooling equilibrium.

The benefit of signalling accrues at the end also in the models of Kremer and Skrzypacz (2007) and Hörner and Vieille (2009), where the signalling action is delaying trade. In the current paper, the sender does not choose whether to (irreversibly) trade or not, but exerts a signalling effort that may vary over time.

One-shot noisy signalling has been studied by Matthews and Mirman (1983), Carlsson and Dasgupta (1997) and Daley and Green (2014). These models describe one-shot interactions (e.g. the seller of a used car offers a warranty to a buyer). The current paper addresses long-term relationships, such as a politician deciding each year how much to cheat on taxes, and voters remembering all past scandals.

Noise interfering with the inference process of the receivers is reminiscent of the signal-jamming literature following Fudenberg and Tirole (1986). In signal-jamming, the incumbent tries to prevent the entrant from learning the entrant's profitability. The present paper describes a situation in which the incumbent tries to convince the entrant that the incumbent is the low-cost type.

Career concerns models (starting with Holmström (1999)) feature noisy effort over time, similarly to repeated signalling models. However, in most of the career concerns literature, the sender does not know his own type and the receivers care about the sender's future actions, not only the type. The present paper focusses on pure signalling, in which the sender knows his type and the action is unproductive—the receivers' utility depends only on the sender's type, not the action.

A variety of reputation models, starting with Kreps and Wilson (1982) and Milgrom and Roberts (1982) share features of dynamic signalling. Costly actions are taken to influence the beliefs of observers, which provide future benefit. An important difference is that most of the reputation literature focusses on private values. That is, the receivers care about the future actions of the sender, not about the type directly. This is precisely the opposite of signalling, where type matters to the receivers, but future actions do not. Furthermore, in this paper, all types are strategic. Most reputation models use commitment types.

Cripps et al. (2004) show that in a wide class of repeated games, reputation for behaviour that is not an equilibrium of the complete information stage

game is temporary and the type must eventually be learned. In the Brownian signalling model of the present paper, both types have positive probability of acquiring a 'false' permanent reputation, in the sense that when signalling stops, belief about the good type may be lower than the prior and belief about the bad type higher. In expectation, beliefs move in the direction of the sender's type, but mistakes have positive probability.

3.2 Extensions

The environment this paper focusses on is pure signalling, in which effort has no direct benefit. A natural question is how the results would change with productive effort. Formal models of productive effort in the frameworks used in this paper are left for future research, but this section discusses some anticipated results.

If the receivers value the signal the effort generates (e.g. work outcomes) in addition to the type, then there is a benefit to signalling even when pooling on no effort is expected. If this benefit is small, the equilibrium set is similar to the case where the benefit is zero. The only change is that signalling can be sustained for a slightly larger set of log likelihood ratios. If the reward the receivers offer the sender for a high signal is large enough, then both types are induced to signal and pooling on positive effort results. The receivers valuing current effort instead of the signal leads to the same conclusions as when the signal is valued directly, provided the effort is unobserved and the signal observed.

If the receivers value the future effort they expect from the sender, then the flow benefit to the sender depends not only on the current log likelihood ratio, but also on the strategy the receivers expect. Suppose the receivers expect higher future effort from the good type than from the bad. Then under good news, the log likelihood ratio drifts down in the absence of a signal. The effort expected from the sender and the expected type fall as the log likelihood ratio falls. The sender then has a lower benefit from a lower log likelihood ratio, so the qualitative properties of the pure signalling model are preserved.

In the bad news case, the future effort expected from either type may fall in the log likelihood ratio, e.g. when pooling is expected at high log likelihood ratios. If the weight the receivers place on future effort is large enough, the flow benefit of the sender may decrease in the log likelihood ratio. This creates an incentive not to signal, so pooling on zero effort is the unique equilibrium. The same effect operates in the Brownian model close to the upper boundary of the signalling region.

Another natural extension is to make the signal depend on the type as well as the effort. A reasonable assumption in that case is that the high type has a higher rate of good signals and the low type a higher rate of bad signals. This leads to a positive marginal benefit of effort even when both types are expected to exert equal effort. If type has a small enough effect on the signal rate, the results that pooling is sometimes the unique equilibrium and pooling on positive effort cannot occur still hold. This is because the marginal cost is positive and, for a small type-dependence of signals, larger than the marginal benefit. A mild

type-dependence of the signal does not change the set of extremal equilibria significantly. If the low type has a zero rate of good news at no effort, pooling still gives the low type the highest payoff among equilibria. Otherwise, the discussion in Section 2.4 applies.

The author conjectures that switched effort equilibria still exist in the bad news model with a slight type-dependence of the signal, because the value function is most likely continuous in this perturbation, just like in the perturbations considered in Sections 2.3 and 2.4. However, closed forms for the value functions or for reasonably tight bounds on the value functions are unlikely to exist, because with any type-dependence of the signal, the HJB equations do not reduce to ODEs. Instead, the substitution of the best responses into the HJB equations results in functional differential-difference equations. It is thus difficult to check whether all the ICs sufficient for a switched effort equilibrium hold. Numerical simulations would have to solve a functional fixed point problem with high accuracy to approximate the value functions and enable checking the ICs.

References

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

Bagwell, K. (1995): "Commitment and observability in games," *Games and Economic Behavior*, 8, 271–280.

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

Cripps, M. W., G. J. Mailath, and L. Samuelson (2004): "Imperfect monitoring and impermanent reputations," *Econometrica*, 72, 407–432.

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.

Dilme, F. (2014): "Dynamic Quality Signaling with Moral Hazard," Working paper 14-019, Penn Institute for Economic Research, Department of Economics, University of Pennsylvania.

Feltovich, N., R. Harbaugh, and T. To (2002): "Too cool for school? Signalling and countersignalling," *RAND Journal of Economics*, 630–649.

Fudenberg, D. and J. Tirole (1986): "A "signal-jamming" theory of predation," *The RAND Journal of Economics*, 17, 366–376.

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

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.

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.

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.

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.

Presman, Ė. L., I. M. Sonin, E. A. Medova-Dempster, and M. A. H. Dempster (1990): *Sequential control with incomplete information: the Bayesian approach to multi-armed bandit problems*, Academic Press.

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.

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

Yushkevich, A. A. (1988): "On the two-armed bandit problem with continuous time parameter and discounted rewards," *Stochastics*, 23, 299–310.