Noisy signalling over time

Sander Heinsalu *

March 7, 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.

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

JEL classification: D82, D83, C73.

---

*Email: sander.heinsalu@yale.edu, website: http://sanderheinsalu.com/ Department of Economics, Yale University. 28 Hillhouse Ave., New Haven, CT 06511, USA.
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 and participants of the 24th International Conference on Game Theory at Stony Brook University, EconCon 2013, the Yale micro theory breakfast and lunch, the Yale Graduate Summer Workshop and the NYU Micro Student Lunch for comments and suggestions. Any remaining errors are the author's. Financial support from Yale University and the Cowles Foundation is gratefully acknowledged.

1 Introduction

Most signalling situations involve noisy observation of the sender's effort. In many, effort is exerted over time. 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.

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. Despite the richness, some results are robust across noise structures. 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. When pooling is not unique, simple informative equilibria exist. 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. Receivers observe a noisy public signal about the effort, rather than the effort itself. The signal is modelled either as a Poisson or a Brownian process. Effort either increases the intensity of the Poisson process (this is called the good news case), or decreases it (bad news). The types only differ in their flow cost of effort. The sender derives a flow benefit directly from the posterior belief of the receivers. Attention is restricted to Markov stationary equilibria.

Some features are common to the Poisson and Brownian models. 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. 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, 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.

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.

Each of the three models has unique features. In the bad news model, the high-cost sender exerts more effort than the low-cost in some equilibria. To describe such equilibria, recall the example of a politician who can exert effort to obey the law. Lawful behaviour decreases the frequency of scandals. The equilibria with higher effort by the corrupt type display four regimes, referred to as early career, insider, scrutiny and tainted. Play starts in the early career, during which the corrupt type exerts positive effort and the honest type 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, the honest type exerts maximum effort and the corrupt type none. Under scrutiny, a scandal leads to a tainted reputation: voters are certain that the politician is corrupt and the politician exerts no effort.

In the good news and Brownian models, 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.

In the Brownian model, belief is bounded away from certainty in all equilibria, unlike in the Poisson models. There need not exist a 'most informative' equilibrium for a given prior. More precisely, for a given equilibrium, define the *signalling region* as the set of beliefs at which the low-cost type of the sender exerts maximal effort and the high-cost type minimal and outside which neither type exerts effort. Two signalling regions corresponding to equilibria at the same parameters might intersect without their union (or any set of beliefs containing the union) being an equilibrium signalling region. The two signalling regions cannot be ranked by informativeness.

Many authors have mentioned the relevance of time (Weiss, 1983; Admati and Perry, 1987) and noise (Matthews and Mirman, 1983) in signalling contexts.

Continuous time signalling with Brownian noise is considered in Daley and Green (2012b), Gryglewicz (2009) and Dilme (2012). 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 (2012b), 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. An overview of Kaya (2009) is given in the online appendix. More distantly related repeated noiseless models are Nöldeke and van Damme (1990) and Swinkels (1999), in which the sender receives the benefit upon stopping signalling.

One-shot noisy signalling in a limit pricing context is studied in Matthews and Mirman (1983). Carlsson and Dasgupta (1997) select equilibria in noiseless one-shot games that are the limits of equilibria in noisy games as noise vanishes. They assume the receiver only has two actions. Daley and Green (2012a) look at perfectly observable effort together with a noisy exogenous signal about the type in a one-shot model. A more detailed discussion of the literature is deferred to Section 4.1.

2 Poisson signalling

This section turns to the main 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 $[0, 1]$ (endowed with the natural Borel σ-algebra) 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 $c_L > c_H > 0$. Effort benefits the sender via its effect on the signal process, which drives the market's log likelihood ratio l. This in turn determines the flow payoff. Before describing this benefit, the signal process, strategies and market expectations must be defined.

The signal is binary, with values in $\{0, 1\}$. The signal 1 occurs at a Poisson rate, and in its absence the signal is 0. In the good news (breakthrough) case, the rate of signal 1 is $e_t \lambda$ at time t. The parameter $\lambda \in (0, \infty)$ is interpreted as the informativeness of effort and e_t denotes the effort at t. The intensity increases in the sender's effort, so the occurrence of the signal is good news about the sender. In the bad news (breakdown) case, the rate of 1 is $(1 - e_t)\lambda$, which decreases in effort. Note that zero effort in the good news case or maximal effort $e = 1$ in the bad news case ensures no signals occur. Given that the values of the signal are fixed, a realization of the signal process is described by the times at which 1 occurs. The receivers observe the public signals, but not the sender's effort. Since the signal is public, the l it leads to is common to all receivers.

A signal sequence is a sequence $(\tau_k)_{k=1}^{\infty}$ of signal times satisfying $0 = \tau_0 < \tau_1 < \tau_2 < \cdots$ and having no accumulation points. The set of signal sequences $\mathcal{H}_\infty$ is endowed with the σ-algebra generated by cylinders. An n-signal public history is a finite sequence $(\tau_1, \ldots, \tau_n, t)$ satisfying $\tau_1 < \cdots < \tau_n < t$, with

[1] Throughout this paper, log likelihood ratio l is used instead of belief $\Pr(H) = \frac{\exp(l)}{1+\exp(l)}$, as formulas simplify significantly in the dynamic models to follow. There is a one-to-one map from log likelihood ratio to belief, so all results can be stated in terms of beliefs.

$t \in (0, \infty)$. The set of n-signal histories is $\mathcal{H}_n$. It inherits the σ-algebra from $\mathcal{H}_\infty$. The set of nonterminal public histories is $\mathcal{H} = \cup_{n \in \mathbb{N} \cup \{0\}} \mathcal{H}_n$. $\mathcal{H}$ inherits a σ-algebra from $\{\mathcal{H}_n\}$ and is a Borel space (Yushkevich, 1980). The terminal public histories are the signal sequences. The truncation of a history $h = (\tau_1, \ldots, \tau_n, t)$ to time $s \leq t$ is $h_s = (\tau_1, \ldots, \tau_m, s)$, with $\tau_m < s$. The notation $\tau_k \in h$ means that under history h, a signal occurs at time τ_k.

A *pure public strategy* is a pair of measurable maps $e = (e_H, e_L)$ from $\mathcal{H}$ into the action set $[0, 1]$. Throughout the paper, only public strategies are considered. In this section, the focus is on pure strategies. Denote the strategy the market expects by $e^* = (e^*_H, e^*_L)$. This notation is also used for equilibrium strategies.

Some sets of histories that are used in defining the updating rule for the log likelihood ratio are defined next. These consist of histories in which a signal occurs at a time at which the strategy expected from the sender is such that the signal is perfectly informative about the type. In the good news case,

$$\begin{aligned}
\mathcal{H}^g_{\max}(e^*) &= \{h \in \mathcal{H} : \exists \tau_k \in h,\ e^*_H(h_{\tau_k}) > 0,\ e^*_L(h_{\tau_k}) = 0\}, \\
\mathcal{H}^g_{\min}(e^*) &= \{h \in \mathcal{H} : \exists \tau_k \in h,\ e^*_H(h_{\tau_k}) = 0,\ e^*_L(h_{\tau_k}) > 0\},
\end{aligned}$$

and in the bad news case,

$$\begin{aligned}
\mathcal{H}^b_{\max}(e^*) &= \{h \in \mathcal{H} : \exists \tau_k \in h,\ e^*_H(h_{\tau_k}) < 1,\ e^*_L(h_{\tau_k}) = 1\}, \\
\mathcal{H}^b_{\min}(e^*) &= \{h \in \mathcal{H} : \exists \tau_k \in h,\ e^*_H(h_{\tau_k}) = 1,\ e^*_L(h_{\tau_k}) < 1\}.
\end{aligned}$$

The updating of the market's log likelihood ratio is described next, starting with the response to signal 1. If the strategy the market expects is e^*, signal 1 occurs at time t, and at t the log likelihood ratio is l_t, then in the good news case the log likelihood ratio jumps to

$$j^g(l_t) = \begin{cases} l_t + \ln\left(\frac{e^*_H(h_t)}{e^*_L(h_t)}\right) \text{ (with } \frac{0}{0} = 1) & \text{if } h_t \notin \mathcal{H}^g_{\max}(e^*) \cup \mathcal{H}^g_{\min}(e^*), \\ l_t & \text{if } h_t \in \mathcal{H}^g_{\max}(e^*) \cup \mathcal{H}^g_{\min}(e^*). \end{cases}$$

In the bad news case the log likelihood ratio jumps to

$$j^b(l_t) = \begin{cases} l_t + \ln\left(\frac{1 - e^*_H(h_t)}{1 - e^*_L(h_t)}\right) \text{ (with } \frac{0}{0} = 1) & \text{if } h_t \notin \mathcal{H}^b_{\max}(e^*) \cup \mathcal{H}^b_{\min}(e^*), \\ l_t & \text{if } h_t \in \mathcal{H}^b_{\max}(e^*) \cup \mathcal{H}^b_{\min}(e^*). \end{cases}$$

Two refinements are built into the $j(l)$ formulas. First, if the efforts expected from the types at the time a signal occurs are such that the signal rate is zero, then by the $\frac{0}{0} = 1$ assumption, the log likelihood ratio does not respond to the signal. Second, signals are ignored when there is certainty about the type (a perfectly informative signal occurred in the past).

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

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

and under bad news, it is

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

where (e_H^*, e_L^*) is the strategy the market expects and $h = (\tau_1, \ldots, \tau_n, t)$ is the history up to t.

Given a signal sequence, 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.

The integrals in (1) and (2) are uniquely defined, because e_L, e_H are bounded and measurable in the σ-algebra of histories, which contains singletons. Fixing a signal sequence, a history is determined by its length t, so e_L, e_H are measurable functions from time to actions.

A *Markov stationary strategy* is a public strategy measurable w.r.t. the log likelihood ratio process. A Markov stationary strategy can be written as a pair of functions $(e_L, e_H) : \overline{\mathbb{R}} \to [0,1]^2$. Subsequently only pure Markov stationary strategies are considered. To simplify the statements to follow, attention is restricted to e_L, e_H piecewise continuous[2] 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.

In this paper, l_0 is treated as a parameter, not a variable, so the strategies are a function of l_0 in addition to l. This is to avoid discussion of log likelihood

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

ratios unreached in equilibrium after any deviation. Define the *reachable set* of log likelihood ratios

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

where l_t is given in Def. 1. The reachable set depends on the strategy the market expects and on l_0. Only behaviour on the reachable set is discussed subsequently. Since no deviation can take l outside the reachable set, behaviour there can be arbitrary. Some $l \in \mathcal{L}(e^*)$ can only be reached by deviating, not by following the equilibrium strategy. For example, L cannot reach $l = \infty$ in the good news case by following $e_L^*(l) = 0\ \forall l \in \mathcal{L}(e^*)$ and H cannot reach $l = -\infty$ in the bad news case by following $e_H^*(l) = 1$ in the interior of the reachable set. To specify behaviour in these cases, assume $e_L^*(\infty) = 0$ and $e_H^*(-\infty) = 0$.

Now that strategies and the log likelihood ratio process have been described, the sender's payoff can be defined. The sender is assumed to derive flow benefit $\beta(l)$ directly from the market's log likelihood ratio l.[3]

The sender's flow utility from effort e and the market's log likelihood ratio l is $\beta(l) - c_\theta e$, where 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 $\beta_{\min}$.

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_θ. The discount rate is $r > 0$. Except for jumps, l evolves deterministically given the market expectations (e_L^*, e_H^*). The jumps occur at Poisson times. Given l at the time of a jump, the size of the jump is deterministic. The expectation in (3) is thus over the jump times of the Poisson signal process induced by $e_\theta(\cdot)$.

---

[3]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 $a^*(l)$.

Since l_0 is a parameter, the payoff starting at l_0 need not in general equal the continuation value from l_0 on when starting at some $\hat{l}_0 \neq l_0$. However, if the strategy the market expects is Markov stationary, then every time l is reached, the continuation value of type θ from l is well defined and is denoted $V_\theta(l)$.[4]

Lemma 1. *$V_H(l) \geq V_L(l)$ $\forall l_0 \in \overline{\mathbb{R}}$ $\forall e^*$ $\forall l \in \mathcal{L}(e^*)$, 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 $l \in \mathbb{R}$.*

All proofs omitted from the text are in Appendix A.

Definition 2. A Markov stationary equilibrium consists of a Markov stationary 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.

The definition implies that on the reachable set, behaviour is optimal from any point on. Therefore the equilibrium concept could also be called Markov perfect. 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 pure equilibria, because the cost and the signal rate are linear in effort. The 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 an equilibrium in which $e^*_L(l) = e^*_H(l)$ $\forall l \in \mathcal{L}(e^*)$.

Lemma 2. *Assume the updating rule satisfies $\forall l \in \mathcal{L}(e^*)$ $\forall t \geq 0$ $\forall s > 0$ if $e^*_H(l) = e^*_L(l)$ and $l_t = l$, then $l_{t+s} = l_t$. Then*

(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 $l_0 \in \overline{\mathbb{R}}$,

[4] The dependence of $V_\theta(l)$ on e^* and l_0 is suppressed in the notation.

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

Proof. If $e_H^*(l) = e_L^*(l)$ and $l_t = l$ imply $l_{t+s} = l_t$, 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. This rules out $e_L^*(l) = e_H^*(l) > 0$ in equilibrium. □

Lemma 2 covers a variety of noise structures, including the Poisson (with the refinements described above) and the Brownian signal structures of this paper. Lemma 2 shows that an extremal equilibrium always exists, because pooling on zero effort is an extremal equilibrium.

Pooling is the unique equilibrium if the benefit of signalling is low enough relative to the cost. It is proved below (Propositions 5 and 9) that if pooling is not the unique equilibrium, then there exists a continuum of nonpooling extremal equilibria.

Proposition 3. *Pooling is the unique equilibrium $\forall l_0 \in \overline{\mathbb{R}}$ if $\frac{\beta_{\max} - \beta_{\min}}{r} \leq \frac{c_H}{\lambda}$.*

The conditions for the existence of an informative equilibrium have the same intuition in the Brownian and the one-shot cases as in the Poisson model: for some initial log likelihood ratio, the benefit to signalling when the receivers expect H to signal and L not to signal has to be high enough to incentivize H to signal. Then an l_0 can be found at which the benefit of signalling is low enough for L not to imitate H, but high enough for H to signal.

Given e^*, the *pooling region* is defined as

$$\mathcal{P}(e^*) = \{l \in \mathcal{L}(e^*) : e_H^*(l) = e_L^*(l) = 0\}.$$

Due to the assumptions that at $l = -\infty$ and ∞, the log likelihood ratio does not respond to signals and $e_L^*(\infty) = e_H^*(-\infty) = 0$, we have $-\infty, \infty \in \mathcal{P}(e^*)\ \forall e^*$.

It is intuitive that when the sender's benefit is concave in the receivers' posterior belief,[5] 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 16 in Appendix A states this formally. The result holds in all the signalling models discussed in this paper.

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, Harbaugh, and To (2002), but the mechanism is quite different. In the bad news model, it is the threat of signalling being required in the future that incentivizes L to signal. This threat is not as severe for H whose signalling cost is lower.

Some effort patterns cannot occur in equilibrium. Lemma 2 ruled out $e_L^*(l) = e_H^*(l) > 0$. Lemma 4 shows that $e_L^*(l) > e_H^*(l)$ cannot occur under some conditions, e.g. when the jump $j(l)$ lands in the pooling region.

Lemma 4. *In equilibrium, the following cannot occur* $\forall l_0 \in \overline{\mathbb{R}}\ \forall l \in \mathcal{L}(e^*)$*:*

(a) $e_L^*(l) > e_H^*(l)$ *and pooling at* $j(l)$*,*

(b) $0 < e_L^*(l) < e_H^*(l) = 1$.

The intuition of the proof of (a) is as follows. The value after a jump is the same for both types, so due to $V_H \geq V_L$, the benefit of the jump is larger (or the loss is smaller) for L than for H, regardless of whether $V_\theta(j(l)) > V_\theta(l)$ or not. The cost of avoiding the jump is larger for L, so it cannot be that L chooses a greater effort to avoid jumps than H. For (b), it is proved that V_L is strictly increasing in the region where $0 < e_L^*(l) < e_H^*(l) = 1$. If L is taking interior effort and H is taking maximal effort, then the jumps go to $l = -\infty$. Given that V_L is strictly increasing, L is indifferent to the jumps

[5] $\beta(l)$ is linear in the posterior belief if it has the form $k_1 \frac{\exp(l)}{1+\exp(l)} + k_2$, with $k_1 > 0$, because the probability corresponding to log likelihood ratio l is $\frac{\exp(l)}{1+\exp(l)}$. The function $k_1 \frac{\exp(l)}{1+\exp(l)} + k_2$ is convex for $l < 0$ and concave for $l > 0$.

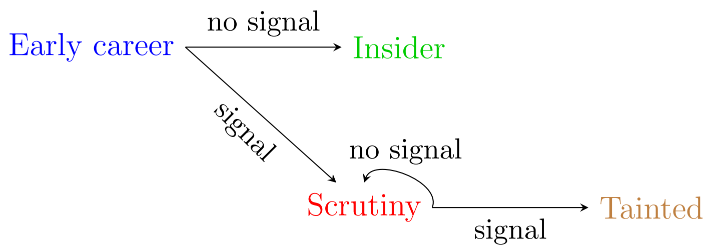

Figure 1: The four regimes in an equilibrium in which $e_L^* > e_H^*$.

at most at a single point, so cannot be taking interior effort over a range of l.

Lemma 4 does not rule out $e_L^* > e_H^*$ occurring in equilibrium, as shown in Example 1. Both $j(l) > l$ and $j(l) < l$ are possible, because if $e_L^*(l) > e_H^*(l)$, then $j(l) > l$. The effort pattern $e_L^* > e_H^*$ 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.

After the example, the set of extremal equilibria is characterized. There exist equilibria in which $e_H^*(l) \in (0,1)$ and $e_L^*(l) = 0$ for some reachable l. This is shown in Lemma 17 in Appendix A. These equilibria can be found numerically, but cannot be solved for in closed form.

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 Proposition 18 in Appendix A provides sufficient conditions for the existence of an equilibrium in which $e_L^* \in (0,1)$ and $e_H^* = 0$ over the interval $(2-\epsilon, l_0]$. In the interval $[2,\infty)$, the efforts are $e_L^* = 0$ and $e_H^* = 1$. Elsewhere, the efforts are zero.

The play in this equilibrium is illustrated in in Figure 1. The initial regime is early career, which in the absence of a signal eventually transitions to insider, but after a signal transitions to scrutiny. Under scrutiny, another signal takes the play to the tainted regime. The incentive for L to exert higher effort 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 is larger for L, so L can be incentivized to positive effort while H takes zero effort. The incentive for initial signalling is provided by the threat of future expectation of signalling.

In the log likelihood ratio space, the four regimes are depicted in Figure 2.

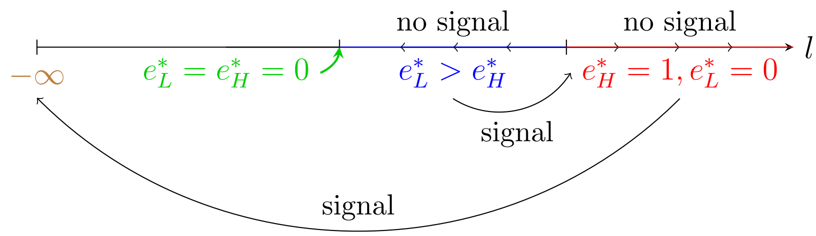

Figure 2: The log likelihood ratios in an equilibrium in which $e_L^* > e_H^*$.

2.1.1 Extremal equilibria

In the class of extremal equilibria, it is w.l.o.g. to consider only the ones in which $e_L^* \equiv 0$ and $e_H^*(l) = 1$ if $l \in [l_0, \bar{l})$ for some $\bar{l} \in \overline{\mathbb{R}}$, with $e_H^*(l) = 0$ otherwise. The interval $[l_0, \bar{l})$ is called the *signalling region*. The reason why extremal equilibria must take an interval form is as follows. 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. This makes L deviate to zero effort. The log likelihood ratio cannot drift across a region on which $e_L^* = e_H^* = 0$ and if $e_H^*(l) = 1$, $e_L^*(l) = 0$, then $j(l) = -\infty$, so there are no jumps into another region where $e_L^* = 0$ and $e_H^* = 1$. The interval $[l_0, \bar{l})$ is assumed open at $\bar{l}$ for three reasons. At $l = \infty$, beliefs do not respond to signals, so both types will choose $e = 0$. Singleton intervals are ruled out by the assumption that e_L^*, e_H^* are piecewise continuous and continuous from the left or the right at jumps. Nonsingleton closed intervals of finite length can be replaced with intervals open on the right without changing the results.

Next, the value functions of the types in extremal equilibria are calculated. After that, bounds on the set of signalling regions are provided, the existence of nonpooling extremal equilibria is proved if the condition in Proposition 3 fails, and finally comparative statics are reported.

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 and a verification theorem (Theorem 4.6 in Presman, Sonin, Medova-Dempster, and Dempster (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) $rV_H(l) = \beta(l) - c_H + \lambda V_H'(l)$ and $rV_L(l) = \beta(l) + \lambda V_L'(l) + \frac{\lambda\beta_{\min}}{r} - \lambda V_L(l)$. 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 after solving the $\bar{l} < \infty$ case). 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.

If $\bar{l} = \infty$, then both types get benefit $\beta(l)$ forever. In addition, L has a flow rate λ of jumps to $\frac{\beta_{\min}}{r}$ and H pays a flow cost c_H forever. The candidate value functions are

$$V_H(l) = \int_l^\infty \frac{\beta(z) - c_H}{\lambda} \exp\left(-r\frac{z-l}{\lambda}\right) dz, \tag{6}$$

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

For every $\epsilon > 0$ there exists $\hat{l} \in \mathbb{R}$ s.t. $\beta_{\max} - \beta(\hat{l}) < \epsilon$, which implies $|\frac{\beta_{\max}-c_H}{r} - V_H(\hat{l})| < \frac{\epsilon}{r}$ and $|\frac{\beta(\hat{l})}{r+\lambda} + \frac{\lambda\beta_{\min}}{r(r+\lambda)} - V_L(\hat{l})| < \frac{\epsilon}{r}$.

Bounds on the set of signalling regions $[l_0, \bar{l})$ for which the incentive constraints (4) are satisfied are provided next. The maximal upper boundary $\max \bar{l}$ that an equilibrium signalling region can have must satisfy the incentive constraint $\lim_{l\to\max\bar{l}-} V_L(l) \leq \frac{\beta_{\min}}{r} + \frac{c_L}{\lambda}$. The limit equals $\frac{\beta(\bar{l})}{r}$ for $\bar{l}$ finite and equals $\frac{\beta_{\max}}{r+\lambda} + \frac{\lambda\beta_{\min}}{r(r+\lambda)}$ for $\bar{l} = \infty$. The benefit to signalling is avoiding the bad signal, so the larger the difference between $\frac{\beta_{\min}}{r}$ and $V_L(l)$, the greater the incentive of L to imitate H. Since V_L is increasing, the L type incentive constraint determines the log likelihood ratio above which signalling cannot continue.

The minimal l_0 at which signalling can start is determined by V_H, which depends on $\bar{l}$. Denote the minimal lower boundary given $\bar{l}$ by $l_0(\bar{l})$. This is finite, because $-\infty$ is in the pooling region by Def. 1. The incentive constraint $\lim_{l\to l_0(\bar{l})} V_H(l) \geq \frac{\beta_{\min}}{r} + \frac{c_H}{\lambda}$ must hold with equality at $l_0(\bar{l})$. Since V_H is strictly increasing, if the IC for H holds at l, then it holds at all $\hat{l} > l$.

There exists an informative extremal equilibrium if the condition in Proposition 3 for pooling to be the unique equilibrium fails.

Proposition 5. *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 $[l_0, l_0+\epsilon)$.*

The intuition for Proposition 5 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 $[l_0, l_0+\epsilon)$.

If there exists one extremal equilibrium with a nonempty signalling region $[l_0, \bar{l})$, then there is a continuum of such equilibria, each corresponding to a particular $\bar{l}' \in (l_0, \bar{l}]$.

Comparative statics are the final item discussed in this section. Welfare is defined as $W(l) = \frac{\exp(l)}{1+\exp(l)} V_H(l) + \frac{1}{1+\exp(l)} V_L(l)$, because the receivers form a competitive market, so their payoff is zero in any equilibrium. 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 6. *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* $\frac{\beta'(\bar{l})}{r} - \frac{c_H}{\lambda} > 0$,

(c) $\frac{dW(l)}{d\bar{l}} > 0$ *iff* $\exp(l)\left[\beta'(\bar{l}) - \frac{c_H r}{\lambda}\right] + \beta'(\bar{l}) - \beta(\bar{l}) + \beta_{\min} > 0$.

Proof. The proof is by taking the appropriate derivatives. □

The condition for $\frac{dV_L(l)}{d\bar{l}} > 0$ holds when $\beta(l) = \left(\frac{\exp(l)}{1+\exp(l)}\right)^n$, $n \in \mathbb{N}$, $n \geq 2$, $\bar{l} < \ln(n-1)$. 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 $l = -\infty$. If $\beta'(\bar{l}) - \beta(\bar{l}) + \beta_{\min} > 0$ as $\bar{l} \to l_0$, then an informative equilibrium yields a higher $V_L(l_0)$ than pooling. If $\beta'(\bar{l}) - \beta(\bar{l}) + \beta_{\min} \leq 0$ for all $\bar{l} > l_0$, then pooling maximizes the payoff of L. This is the case if $\beta(l) = \frac{\exp(l)}{1+\exp(l)}$ (the benefit from the receivers' belief is linear in the belief).

If the condition for $\frac{dV_H(l)}{d\bar{l}} > 0$ holds as $\bar{l} \to l_0$, then an informative equilibrium gives H a higher payoff than pooling. If the condition fails at all $\bar{l} > l_0$, then pooling maximizes the payoff of H. The interpretation of the condition 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.

The effect of increasing $\bar{l}$ on welfare is a combination of the effects on H and L, with a weight $\exp(l)$ on the payoff of H and a weight 1 on L.

2.2 The good news case

The results about the good news model are presented next. The L type always prefers the pooling equilibrium under good news, even when the flow benefit from the receivers' log likelihood ratio is convex. This is reminiscent of L preferring pooling to any separating equilibrium in noiseless models, but differs from the other noisy models considered in this paper.

Some preliminary observations about the equilibrium efforts are collected in the following lemma. Since the lemma rules out $e_L^* > e_H^*$, the log likelihood ratio can only jump up: $j(l) \geq l$ in the good news case.

Lemma 7. *$\forall l_0 \in \overline{\mathbb{R}}\ \forall l \in \mathcal{L}(e^*)$, 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 $e_H^*(l) = 1$.*

Some equilibria with $e_L^*(l), e_H^*(l) \in (0,1)$ at some $l \in \mathcal{L}(e^*)$ can also be ruled out, as shown in Lemma 19 in Appendix A.

Restricting attention to extremal equilibria,[6] Lemma 7 implies that in equilibrium the only possible effort combinations are $e_H^* = e_L^* = 0$ and $e_H^* = 1$, $e_L^* = 0$. An extremal equilibrium must therefore be an interval of log likelihood ratios $(\underline{l}, l_0]$ on which $e_H^* = 1$, $e_L^* = 0$ and outside which $e_H^* = e_L^* = 0$.[7] The lower boundary of the interval can be infinite. The upper boundary is finite, as shown in Lemma 8.

Lemma 8. *$\forall l_0 \in \overline{\mathbb{R}}\ \exists \hat{l} \in \mathbb{R}$ s.t. in any equilibrium $\forall l \in \mathcal{L}(e^*) \cap [\hat{l}, \infty]$, $e_H^*(l) = e_L^*(l) = 0$.*

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. Comparative statics are derived at the end of the section.

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

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

[6] Focussing on extremal equilibria is a restriction. There exist equilibria where $e_L^* \in (0,1)$ and $e_H^* = 1$, as shown in Lemma 20 in Appendix A.

[7] It is w.l.o.g. to consider only one interval, because l cannot drift across a region where $e_H^* = e_L^* = 0$ and if $e_H^* = 1$, $e_L^* = 0$, then jumps are to $l = \infty$. The interval is open at $\underline{l}$, because e_H^*, e_L^* are piecewise continuous and continuous from the left or the right at jumps. This rules out singleton intervals. The interval $[-\infty, l_0]$ cannot be an equilibrium, because beliefs do not respond to signals at $l = -\infty$, so both types will choose zero effort there.

In the pooling region, 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 $V_\theta(l)$. The probability increases with 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 $e^*_H=1$ and $e^*_L=0$ into the HJB equations of H and L. The HJB equations become the ODEs $\lambda V_H'(l)+(\lambda+r)V_H(l)=\beta(l)+\frac{\lambda\beta_{\max}}{r}-c_H$ and $\lambda V_L'(l)+rV_L(l)=\beta(l)$. 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 after solving the $\underline{l}>-\infty$ case.

The solutions of the ODEs are

$$\begin{aligned}V_H(l)&=\exp\left(-(r+\lambda)\,\frac{l-\underline{l}}{\lambda}\right)\frac{\beta(\underline{l})}{r}\\&\quad+\int_{\underline{l}}^{l}\left[\frac{\beta(z)-c_H}{\lambda}+\frac{\beta_{\max}}{r}\right]\exp\left(-(r+\lambda)\,\frac{l-z}{\lambda}\right)dz,\qquad(7)\\V_L(l)&=\exp\left(-r\frac{l-\underline{l}}{\lambda}\right)\frac{\beta(\underline{l})}{r}+\int_{\underline{l}}^{l}\frac{\beta(z)}{\lambda}\exp\left(-r\frac{l-z}{\lambda}\right)dz.\end{aligned}$$

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.

Next, the signalling regions with $\underline{l}=-\infty$ are discussed. Type L has no chance of a jump and gets the discounted flow benefit forever, so L's candidate value is

$$V_L(l) = \int_{-\infty}^{l} \frac{\beta(z)}{\lambda} \exp\left(-r\frac{l-z}{\lambda}\right) dz.$$

Since β is bounded, for any $\epsilon > 0$ there exists $\hat{l} \in \mathbb{R}$ s.t. $\beta(\hat{l}) - \beta_{\min} < \epsilon$, which implies $V_L(\hat{l}) - \frac{\beta_{\min}}{r} < \frac{\epsilon}{r}$. Therefore $\lim_{l \to -\infty} V_L(l) = \frac{\beta_{\min}}{r}$. The term $\beta(\underline{l}) = \beta_{\min}$ does not appear in the $V_L(l)$ expression, because it takes an infinite time to reach $\underline{l} = -\infty$ and, as can be seen from (7), the discounting then makes the (finite) $\beta(\underline{l})$ vanish.

Type H has a constant rate λ of jumps to $l = \infty$ and pays a flow cost c_H forever, in addition to getting the flow benefit. The candidate value of H is

$$V_H(l) = \int_{-\infty}^{l} \left[\frac{\beta(z) - c_H}{\lambda} + \frac{\beta_{\max}}{r}\right] \exp\left(-(r+\lambda)\frac{l-z}{\lambda}\right) dz.$$

The limiting value is $\lim_{l \to -\infty} V_H(l) = \frac{\beta_{\min} - c_H}{r+\lambda} + \frac{\lambda \beta_{\max}}{r(r+\lambda)}$. The RHS is greater than $\frac{\beta_{\min}}{r}$ iff $\frac{\beta_{\max} - \beta_{\min}}{r} \geq \frac{c_H}{\lambda}$.

The signalling regions $(\underline{l}, l_0]$ for which the incentive constraints are satisfied can now be characterized. Due to $V_L' > 0$, if $V_L(l) \geq \frac{\beta_{\max}}{r} - \frac{c_L}{\lambda}$ holds at l, then it holds for all $\hat{l} > l$. If L is deterred from imitating H at l, then L is also deterred at all $\hat{l} > l$, because the benefit to imitation is the difference between the payoff of being believed to be the H type and the current value. The higher l, the higher the current value, so the lower the incentive to exert effort. Since l drifts down in the absence of a signal, the L type incentive constraint determines the log likelihood ratio below which signalling must stop. The minimal lower boundary $\min \underline{l} \geq -\infty$ that an equilibrium signalling region can have must satisfy the incentive constraint $\lim_{l \to \min \underline{l}+} V_L(l) \geq \frac{\beta_{\max}}{r} - \frac{c_L}{\lambda}$, where $\lim_{l \to \min \underline{l}+} V_L(l) = \frac{\beta(\underline{l})}{r}$ and $\beta(-\infty) = \beta_{\min}$. If the expectation was for signalling to continue at $\min \underline{l}$, then L would deviate to $e = 1$.

The maximal initial log likelihood ratio at which signalling can start is determined by V_H, which depends on $\underline{l}$. Due to $V_H' > 0$, if the incentive constraint $V_H(l) \leq \frac{\beta_{\max}}{r} - \frac{c_H}{\lambda}$ holds at l, then for all $\hat{l} < l$ we have $V_H(\hat{l}) < \frac{\beta_{\max}}{r} - \frac{c_H}{\lambda}$. If H is incentivized to signal at l, then H is also incentivized at all $\hat{l} < l$, because the benefit to signalling is the difference between the payoff of being believed to be the H type and the current value. Since l drifts down in the absence of a signal, the H type incentive constraint determines the log likelihood ratio $l_0(\underline{l})$ above which signalling cannot start. If the expectation was for signalling to start above $l_0(\underline{l})$, then H would deviate to $e = 0$. By Lemma 8, $l_0(\underline{l}) < \infty$ for any $\underline{l}$.

The parameter values for which there exists an informative extremal equilibrium are given in Proposition 9.

Proposition 9. *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* $[l_0, l_0+\epsilon)$*.*

If there exists one extremal equilibrium with a nonempty signalling region $(\underline{l}, l_0]$, then there is a continuum of such equilibria, each corresponding to a particular $\underline{l}' \in [\underline{l}, l_0)$.

Comparative statics of extremal equilibria are presented next. Based on (7), $V_H(l)$ 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 for comparative statics. 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 10. *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* $\beta'(\underline{l}) - \beta_{\max} + \beta(\underline{l}) + \frac{c_H r}{\lambda} > 0$*,*

(c) $\frac{dW(l)}{d\underline{l}} > 0$ *iff* $\exp(l)\left[\beta'(\underline{l}) - \beta_{\max} + \beta(\underline{l}) + \frac{c_H r}{\lambda}\right] + \beta'(\underline{l}) > 0$*.*

Proof. The proof is by taking the relevant derivatives. □

Comparing extremal equilibria, $\frac{dV_L(l)}{d\underline{l}} > 0$, so pooling always gives L the highest payoff. This holds 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.

The result that pooling always gives L the highest payoff in the good news model holds not just for extremal equilibria. By Lemma 7, 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.

If the condition for $\frac{dV_H(l)}{dl} > 0$ holds at every $\underline{l} < l_0$, then pooling gives H the highest payoff. If the condition fails as $\underline{l} \to l_0$, then an informative equilibrium gives H a higher payoff than pooling. The condition 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 condition for welfare to increase in $\underline{l}$ holds when $\beta(l) = \frac{\exp(l)}{1+\exp(l)}$, so pooling gives the highest welfare when the sender's benefit from the receivers' belief is linear. The condition for welfare to increase in $\underline{l}$ is a combination of the effects on the payoffs of H and L. The initial $\exp(l)$ multiplier is a weight on the payoff of H. The weight increases in the probability that the sender's type is H. The term in the square brackets is the effect of increasing $\underline{l}$ on the payoff of H, which is discussed above. The final term $\beta'(\underline{l})$ (with a weight of 1) is positive and describes L's benefit from an increase in $\underline{l}$, namely that the payoff upon reaching $\underline{l}$ is larger.

2.3 Final remarks on the Poisson model

Some of the restrictive assumptions in the Poisson games considered above are relaxed in the online appendix. Including both good and bad news in the game, with good signals occurring at rate $\lambda_g e$ and bad signals at rate $\lambda_b(1-e)$, the set of extremal equilibria is similar to the good news case when $\lambda_g > \lambda_b$ and similar to the bad news case when $\lambda_g < \lambda_b$. If $\lambda_g = \lambda_b$, then the log likelihood ratio stays constant in the absence of signals and jumps when a signal occurs. Adding a small positive rate of bad news to the good news model does not change the result that L always prefers pooling. The set of extremal effort equilibria remains similar to the case with only good news. A small positive rate of good news in the bad news model does not significantly affect extremal equilibria.

The assumption that zero effort in the good news case and maximal effort in the bad news case ensure the absence of signals is rather stark. However, as shown in the online appendix, a low Poisson intensity $\epsilon > 0$ of signals even with zero effort in the good news case or with maximal effort in the bad news case does not qualitatively change the set of extremal equilibria.

The positive lower bound on the signal rate in the good news model does overturn the result that L always prefers pooling.

3 Signalling with Brownian noise

The Poisson signalling game corresponds to an environment where information is revealed by rare and significant events. There are situations in which a gradual and continuous information revelation is more realistic. This section turns to the gradual information revelation case and models the signal process as a Brownian motion. In this model, the union of equilibrium signalling regions need not be an equilibrium signalling region. The union would mean the receivers expect 'too much' signalling, which induces one of the types to deviate. An analogous result holds in the one-shot noisy signalling model in the online appendix, where for some parameter values an equilibrium with $e_L^* = 0$, $e_H^* = 1$ exists iff the initial log likelihood ratio belongs to one of two disjoint finite intervals. The expectation of 'too much' signalling between these intervals induces L to imitate.

The setup in the Brownian case is similar to the Poisson model. The sender's effort process $(e_t)_{t\geq 0}$ now controls the drift of a signal process $(X_t)_{t\geq 0}$ given by

$$dX_t = e_t dt + \sigma dB_t,$$

where $X_0 \in \mathbb{R}$ is a given parameter, B_t is standard Brownian motion and $\sigma > 0$. Denote the filtration generated by $(X_t)_{t\geq 0}$ by $(\mathcal{F}_t)_{t\geq 0}$. The receivers at time t observe $(X_\tau)_{\tau\in[0,t]}$, but not the sender's type or present or past actions. Based on the signal, the receivers update their log likelihood ratio. The log likelihood ratio process $(l_t)_{t\geq 0}$ is adapted to $(\mathcal{F}_t)_{t\geq 0}$.

The flow utility of a sender of type θ from action $e \in [0,1]$ and the receivers' log likelihood ratio l is $\beta(l) - c_\theta e$. Assume β is bounded, Lipschitz, strictly increasing and twice continuously differentiable on $\overline{\mathbb{R}}$.

A *pure public strategy* of the sender is a pair of random processes (e_L, e_H), each taking values in $[0,1]$ and adapted to $(\mathcal{F}_t)_{t\geq 0}$. A *Markov stationary strategy* of the sender is a pure public strategy measurable w.r.t. the log likelihood ratio process $(l_t)_{t\geq 0}$. It can be written as a pair of measurable functions $(e_L, e_H) : \overline{\mathbb{R}} \to [0,1]^2$. The state variable is the receivers' log likelihood ratio l. Henceforth, strategy is understood as a Markov stationary strategy.

Given the initial log likelihood ratio $l_0 \in \mathbb{R}$ and the strategy (e_L^*, e_H^*) expected from the sender, the receivers update the log likelihood ratio[8]

$$dl_t = \sigma^{-2}(e_H^*(l_t) - e_L^*(l_t))[dX_t - \frac{1}{2}e_H^*(l_t)dt - \frac{1}{2}e_L^*(l_t)dt]. \quad (8)$$

The initial log likelihood ratio l_0 is a parameter, so a strategy is a function of l_0 in addition to l. The *reachable set* of log likelihood ratios is

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

where l_t is defined in (8). Only behaviour on the reachable set is discussed subsequently. Behaviour outside $\mathcal{L}(e^*)$ can be arbitrary, because no deviation can take l there.

Definition 3. A Markov stationary equilibrium consists of a strategy (e_H^*, e_L^*) and a log likelihood ratio process $(l_t)_{t\geq 0}$ s.t.

1. given $(l_t)_{t\geq 0}$, e_θ^* solves

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

   where the expectation is over the process $(l_t)_{t\geq 0}$,

2. given (e_H^*, e_L^*) and l_0, the process $(l_t)_{t\geq 0}$ is derived from Bayes' rule (8).

Given e^*, the *pooling region* is

$$\mathcal{P}(e^*) = \{ l \in \mathcal{L}(e^*) : e_H^*(l) = e_L^*(l) \}.$$

The complement of the pooling region in $\mathcal{L}(e^*)$ is called the *signalling region.* The *pooling equilibrium* where $e_H^* = e_L^* \equiv 0$ always exists by Lemma 2. A nonpooling equilibrium exists if the sender is patient enough. The proof is postponed to Proposition 15.

Since l_0 is a parameter, the payoff starting at l_0 need not in general equal the continuation value from l_0 after starting at $\hat{l} \neq l_0$. However, if the strategy the market expects is Markov stationary, then every time l is reached, the continuation value $V_\theta(l)$ of type θ from l on is well defined. Some observations about the value functions are formalized in the following lemma.

[8]The updating rule for the log likelihood ratio is derived from the continuous time Bayes' rule for probability (Liptser and Shiryaev (1977) Theorem 9.1) using Itō's formula.

Lemma 11. $V_H(l) \geq V_L(l)$ $\forall l_0 \in \overline{\mathbb{R}}$ $\forall e^*$ $\forall l \in \mathcal{L}(e^*)$, *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* $l \in \mathbb{R}$. V_θ *is strictly increasing.*

The results in Lemma 11 hold more generally. The fact that $V_H(l) \geq V_L(l)$ is independent of the signal structure and the cost function—it only needs that each effort level costs less for H than for L. The value functions are bounded independently of the signals and costs. The strict monotonicity of V_θ does not carry over to the Poisson case, but it holds whenever the paths of the l process are continuous for any strategy expected or chosen.

Equilibrium strategies are partially characterized in the following lemma. If L is expected to take higher effort than H, then a higher signal lowers the log likelihood ratio. With increasing V_θ this implies both types optimally choose 0. This result holds whenever the signal is such that the paths of the l process are continuous.

Lemma 12. *Fix* $l_0 \in \overline{\mathbb{R}}$. *In any equilibrium,* $\nexists l \in \mathcal{L}(e^*)$ *satisfying* $e_L^*(l) = e_H^*(l) > 0$ *or* $e_L^*(l) > e_H^*(l)$.

It is clear from (8) that once the log likelihood ratio process reaches the pooling region, l stays constant forever, so signalling must occur in an interval of l containing l_0. Starting inside the interval, if the l process reaches the boundary, then in the next instant it enters the pooling region due to the rapidly varying Brownian motion driving l. For this reason, it is w.l.o.g. to consider only open signalling regions. In light of this and Lemma 12, it is w.l.o.g. to consider only equilibria in which outside an interval of log likelihood ratios $(\underline{l}, \overline{l})$, both types choose action 0 and inside that interval, $e_L^* < e_H^*$.

The subsequent focus of this section is on extremal equilibria. It can be shown that if effort only takes values zero and one, then outside an interval of log likelihood ratios $(\underline{l}, \overline{l}) \ni l_0$, both types choose action 0 and inside that interval, $e_L^* = 0$, $e_H^* = 1$. In $(\underline{l}, \overline{l})$, the l process is a simple Brownian motion with drift either $\frac{1}{2}$ or $-\frac{1}{2}$, depending on whether the sender's chosen action is $e = 1$ or 0.

The value functions are calculated next, using the HJB equations and a verification theorem. Payoff comparisons readily apparent from the value functions are noted while deriving the value functions. Necessary conditions for equilibrium are derived. After that, the existence of informative extremal equilibria is proved, followed by comparative statics, which are mostly numerical.

It is clear that for $\hat{\beta}(\mu) = \beta(\ln \frac{\mu}{1-\mu})$ concave in μ (the receivers' belief), L prefers pooling to any other equilibrium, because the expected posterior is lower than the prior and the variance in the posterior does not increase the payoff. The concavity of $\hat{\beta}$ in μ implies the condition in Proposition 13 that is sufficient for L to prefer pooling. If pooling gives the highest payoff to L, then in the class of extremal equilibria, L's payoff is higher in an equilibrium with a smaller signalling region.

Proposition 13. *If $\beta(\ln(z))$ is concave in z, then V_L is higher in pooling than in any extremal equilibrium. In that case if V_{L2} is the value of L in an equilibrium with signalling region $(\underline{l}_2, \bar{l}_2)$ and V_{L1} is the value of L in an equilibrium with signalling region $(\underline{l}_1, \bar{l}_1) \subset (\underline{l}_2, \bar{l}_2)$, then $V_{L1}(l) \geq V_{L2}(l)$ for all $l \in (\underline{l}_1, \bar{l}_1)$. If $\beta(\ln(z))$ is convex in z, then V_L is lower in pooling than in any extremal equilibrium and $V_{L1}(l) \leq V_{L2}(l)$ for all $l \in (\underline{l}_1, \bar{l}_1)$.*

The idea of the proof is to transform the l process into a zero-drift process $f(l)$ by an increasing transformation f using Itō's lemma. The benefit function β is simultaneously transformed by the inverse of f. Pooling gives L a higher payoff than an informative extremal equilibrium iff the transformed benefit function is concave in $f(l)$. It turns out $f(l)$ is the likelihood ratio $\exp(l)$.

An example where L prefers pooling has $\beta(l) = \left(\frac{\exp(l)}{1+\exp(l)}\right)^n$, $n \geq 2$ and $\bar{l} \leq \ln \frac{n-1}{2}$.

Next, the HJB equations are solved and a verification theorem is used to check that the solutions of the HJB equations coincide with the value functions. The HJB equation of type θ is

$$rV_\theta(l) = \beta(l) + \max\left\{-c_\theta e + V_\theta'(l)\sigma^{-2}\left(e - \frac{1}{2}\right)\right\} + \frac{1}{2}V_\theta''(l)\sigma^{-2}.$$

Given the signalling region $(\underline{l}, \bar{l})$ the receivers expect, the optimal strategy of type θ is to choose

$$e_\theta(l) = \begin{cases} \mathbf{1}\left\{-c_\theta + V_\theta'(l)\sigma^{-2} \geq 0\right\} & \text{if } l \in (\underline{l}, \bar{l}), \\ 0 & \text{if } l \notin (\underline{l}, \bar{l}). \end{cases}$$

The incentive constraints for H to choose $e_H(l) = 1$ and L to choose $e_L(l) = 0$ in the signalling region are

$$V_H'(l) \geq c_H \sigma^2, \qquad V_L'(l) \leq c_L \sigma^2. \tag{9}$$

Call these IC_H and IC_L. After finding the candidate equilibrium, it must be verified that the ICs hold at every point in the signalling region.

Set the chosen actions equal to the equilibrium actions. The HJB equations become the pair of linear second-order ODEs

$$\begin{aligned} rV_H(l) &= \beta(l) - c_H + \frac{1}{2}V_H'(l)\sigma^{-2} + \frac{1}{2}V_H''(l)\sigma^{-2}, \\ rV_L(l) &= \beta(l) - \frac{1}{2}V_L'(l)\sigma^{-2} + \frac{1}{2}V_L''(l)\sigma^{-2}. \end{aligned}$$

This is where using the log likelihood ratio instead of the belief is helpful—with belief, the ODEs do not have constant coefficients. After solving the ODEs for V_L, V_H, the ICs as well as the smoothness conditions for the verification theorem must be checked at every point in the signalling region.

The solutions V_θ to the ODEs are the sum of the general solution $C_{\theta 1} y_{\theta 1} + C_{\theta 2} y_{\theta 2}$ of the homogeneous equation and a particular solution $y_{\theta p}$ of the inhomogeneous equation. The general solutions for H and L respectively are

$$\begin{aligned} &C_{H1} \exp\left(l\frac{-1-\sqrt{1+8r\sigma^2}}{2}\right) + C_{H2} \exp\left(l\frac{-1+\sqrt{1+8r\sigma^2}}{2}\right), \\ &C_{L1} \exp\left(l\frac{1-\sqrt{1+8r\sigma^2}}{2}\right) + C_{L2} \exp\left(l\frac{1+\sqrt{1+8r\sigma^2}}{2}\right). \end{aligned}$$

Using d'Alembert's method, the particular solutions are

$$\begin{aligned} y_{Hp} = &-\frac{c_H}{r} \\ &+ \frac{2\sigma^2}{\sqrt{1+8r\sigma^2}} \exp\left(l\frac{-1-\sqrt{1+8r\sigma^2}}{2}\right) \int \beta(l) \exp\left(l\frac{1+\sqrt{1+8r\sigma^2}}{2}\right) dl \\ &- \frac{2\sigma^2}{\sqrt{1+8r\sigma^2}} \exp\left(l\frac{-1+\sqrt{1+8r\sigma^2}}{2}\right) \int \beta(l) \exp\left(l\frac{1-\sqrt{1+8r\sigma^2}}{2}\right) dl, \end{aligned}$$

$y_{Lp} =$

$$\frac{2\sigma^2}{\sqrt{1+8r\sigma^2}} \exp\left(l\frac{1-\sqrt{1+8r\sigma^2}}{2}\right) \int \beta(l) \exp\left(l\frac{-1+\sqrt{1+8r\sigma^2}}{2}\right) dl$$
$$- \frac{2\sigma^2}{\sqrt{1+8r\sigma^2}} \exp\left(l\frac{1+\sqrt{1+8r\sigma^2}}{2}\right) \int \beta(l) \exp\left(l\frac{-1-\sqrt{1+8r\sigma^2}}{2}\right) dl,$$

where the integrals are nonelementary even for simple functional forms of β, e.g. for $\beta(l) = \frac{\exp(l)}{1+\exp(l)}$, which describes linear benefit from the receivers' belief.

Imposing the boundary conditions $V_\theta(\underline{l}) = \frac{\beta(\underline{l})}{r}$ and $V_\theta(\bar{l}) = \frac{\beta(\bar{l})}{r}$, the constants in the general solution for H are

$$C_{H1} = \frac{y_{H2}(\bar{l})[\frac{\beta(\underline{l})}{r} - y_{Hp}(\underline{l})] - y_{H2}(\underline{l})[\frac{\beta(\bar{l})}{r} - y_{Hp}(\bar{l})]}{y_{H1}(\underline{l})y_{H2}(\bar{l}) - y_{H2}(\underline{l})y_{H1}(\bar{l})},$$
$$C_{H2} = \frac{-y_{H1}(\bar{l})[\frac{\beta(\underline{l})}{r} - y_{Hp}(\underline{l})] + y_{H1}(\underline{l})[\frac{\beta(\bar{l})}{r} - y_{Hp}(\bar{l})]}{y_{H1}(\underline{l})y_{H2}(\bar{l}) - y_{H2}(\underline{l})y_{H1}(\bar{l})}.$$

The constants for L are determined by a similar expression, replacing the H subscripts with L.

Now that all components of the solutions of the HJB equations have been found, it can be verified that they coincide with the candidate value functions. The ICs remain to be checked.

Lemma 14. *The solutions V_H, V_L of the HJB equations equal the candidate value functions in the signalling region. The Markov controls for the HJB equations maximize the candidate value functions.*

Based on the candidate value function expressions, the signalling region depends on r and σ^2 only through their product $r\sigma^2$. Closed form comparative statics results are not available for parameters other than c_L due to the complexity of the V_θ expressions. Numerical simulations will be used instead. As to c_L, the LHS of IC_L in (9) does not contain c_L, so there exists $\hat{c}_L$ s.t. for $c_L < \hat{c}_L$, IC_L fails and for $c_L \geq \hat{c}_L$, IC_L holds.

Before turning to numerics, the conditions for the existence of nontrivial extremal equilibria are provided. The intuition of the conditions is similar to the Poisson case—an informative equilibrium exists iff the maximal benefit from signalling is large enough to incentivize H to signal. To see the similarity, consider σ^2 analogous to $\frac{1}{\lambda}$ in the Poisson models.

Proposition 15. *If $\exists l \in \mathbb{R}$ with $\frac{\beta'(l)}{r} > c_H\sigma^2$, then $\exists l_0 \in \mathbb{R}$ contained in the signalling region of an extremal equilibrium. If $\nexists l \in \mathbb{R}$ with $\frac{\beta'(l)}{r} \geq c_H\sigma^2$, then pooling is the unique equilibrium.*

The idea of the proof is that the slopes of the value functions are close to the slope of β for small signalling intervals. If the slope of the benefit function is high enough at some point l to incentivize H to signal at l (provided the receivers expect $e_L^*(l) = 0$, $e_H^*(l) = 1$), then there exists another point l_0 with $\beta'(l_0)$ low enough not to incentivize L to signal, but high enough to still incentivize H. An informative extremal equilibrium can then be constructed with a signalling interval containing l_0.

It is clear that for any $l \in \mathbb{R}$, $\sigma^2 > 0$, $c_L > c_H > 0$ and strictly increasing $\beta(l)$, there exists $r \in (0, \infty)$ that makes the sufficient condition in Proposition 15 hold. One can always find a level of patience for a nontrivial extremal equilibrium to exist.

Numerical comparative statics on the set of signalling regions are presented next. As in the Poisson signalling game, for some initial log likelihood ratios there is a continuum of informative equilibria. Until the end of this section, it is assumed that $\beta(l) = \frac{\exp(l)}{1+\exp(l)}$, so the sender's benefit from the receivers' belief equals the belief. The figures to follow depict signalling intervals $(\underline{l}, \bar{l})$ as points on a plane, with the x-coordinate of the point equalling $\underline{l}$ and the y-coordinate equalling $\bar{l}$.

For $c_H = 0.1$, $c_L = 0.24$ and $r = \sigma^2 = 1$, the region where the ICs hold is depicted in panel (c) of Figure 3 as the shaded area. Panel (a) shows the area where IC_H holds and panel (b) the area where IC_L holds. The shaded area on panel (c) is the intersection of panels (a) and (b). The disconnectedness of the set of equilibrium signalling regions is in part due to restricting attention to extremal equilibria.

The effect of increased patience or reduced noise on the ICs is shown in Figure 4, where $c_H = 0.1$, $c_L = 0.24$, $r = 1$ and $\sigma^2 = 0.5$. Note the different scale of the axes compared to Figure 3. Since r and σ^2 affect the ICs only through their product, reducing σ^2 by half has the same effect as reducing r by half.

There need not exist a signalling region containing all others. Such a signalling region is the point at the upper left corner of the shaded area of panel (c), i.e. a point that is simultaneously at maximal horizontal and vertical distance from the diagonal. Figure 5 shows that for $c_H = 0.15$, $c_L = 0.28$ and $r = \sigma^2 = 1$, a higher $\underline{l}$ permits a higher $\bar{l}$ for a signalling

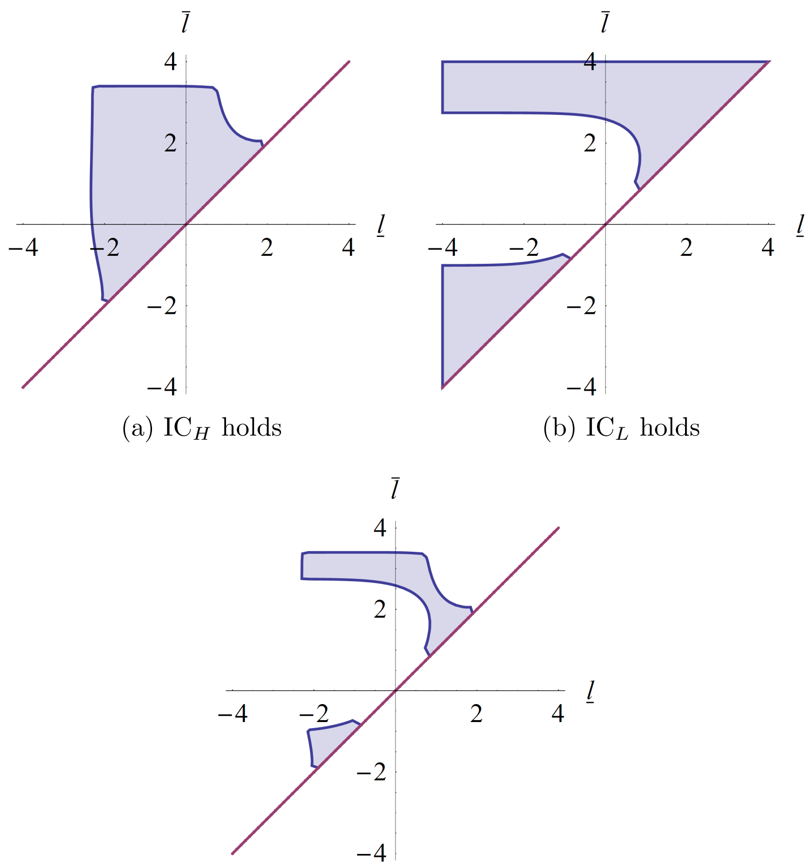

(a) IC_H holds

(b) IC_L holds

(c) Both ICs hold

Figure 3: Region where ICs hold (shaded) for $c_H = 0.1$, $c_L = 0.24$ and $r = \sigma^2 = 1$.

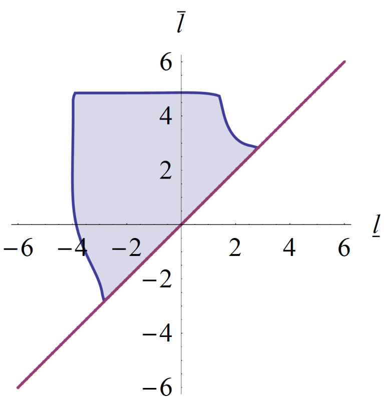

(a) IC_H holds

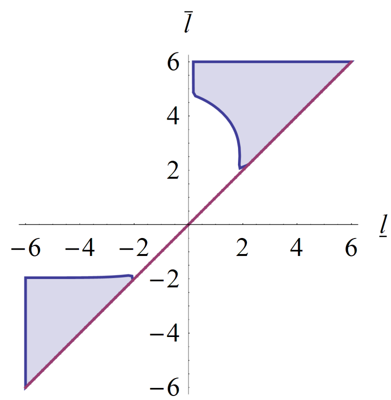

(b) IC_L holds

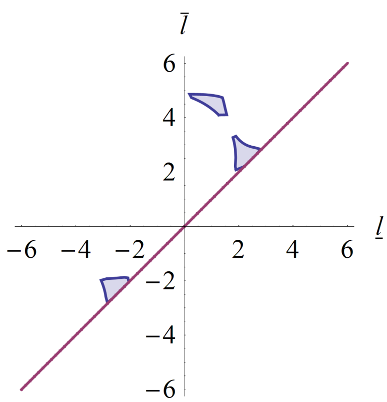

(c) Both ICs hold

Figure 4: Region where ICs hold (shaded) for $c_H = 0.1$, $c_L = 0.24$, $r = 1$ and $\sigma^2 = 0.5$.

region. Therefore the union of two equilibrium signalling regions need not be an equilibrium signalling region. This distinguishes the game with Brownian noise from the Poisson signalling game, the repeated noiseless game and the one-shot noisy and noiseless games.

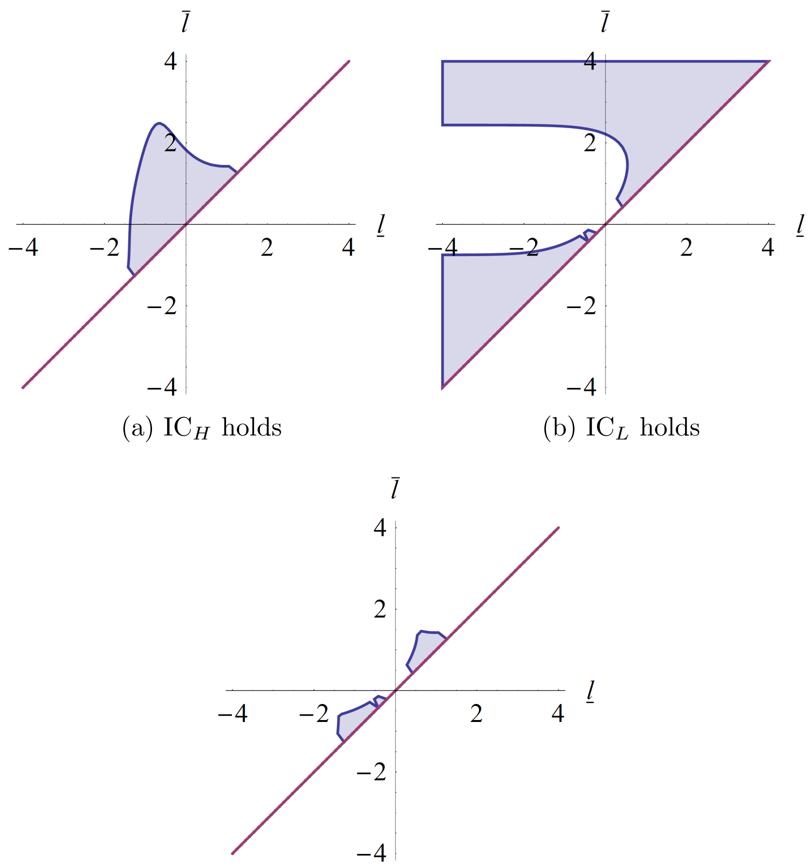

(a) IC_H holds

(b) IC_L holds

(c) Both ICs hold

Figure 5: Region where ICs hold (shaded) for $c_H = 0.15$, $c_L = 0.28$, $r = 1$ and $\sigma^2 = 1$.

In a given informative equilibrium, the H type payoff can be higher or lower than the pooling payoff $\frac{\exp(l)}{r(1+\exp(l))}$ for different log likelihood ratios. This is illustrated in Figure 6, where V_H is strictly higher than $\frac{\exp(l)}{r(1+\exp(l))}$ for $l \in (-1.5, -0.2)$ and strictly lower for $l \in (-0.2, 3)$. In this equilibrium, the comparison of informative equilibrium and pooling payoffs of H accords well with Spence (1973), where for a higher fraction of H in the population, the payoff difference (separating minus pooling) for H is lower. In Spence's model, the reason is that for a higher l_0 there is less scope for the log likelihood ratio to rise ($l = \infty$ after the high action). In the present model this mechanism does not work, because the rise in belief after a good signal is highest for intermediate l_0. Correspondingly the Spence intuition does not always hold in the continuous time model. Close to the upper bound of the signalling region, the payoff from the informative equilibrium rises to the pooling one as l_0 increases, so in that region, the informative equilibrium payoff minus the pooling payoff rises in l.

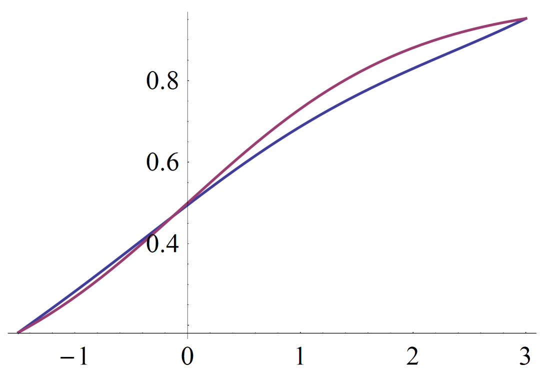

Figure 6: V_H for signalling region $(-1.5, 3)$ (the curve that is lower on the right), and $\frac{\exp(l)}{r(1+\exp(l))}$. The parameters are $c_H = 0.1$ and $r = \sigma^2 = 1$.

For the signalling region $(0, 3)$, with $c_H = 0.1$ and $r = \sigma^2 = 1$, the informative equilibrium payoff of H is below pooling in the whole signalling region. Close to the upper bound, the informative equilibrium payoff minus the pooling payoff rises in l, while close to the lower bound, it falls in l. This pattern is reversed in the informative equilibrium with signalling region $(-3, 0)$, $c_H = 0.1$ and $r = \sigma^2 = 1$. In that case, the informative equilibrium payoff is above pooling in the whole signalling region. Within the signalling region of a given informative equilibrium, there is always a region where the payoff difference with pooling moves in the opposite direction to the prediction of Spence (1973). Comparing equilibria with $\bar{l} \leq 0$ to those with $\underline{l} \geq 0$, the Spence pattern holds—shifting the signalling region up raises the informative equilibrium payoff minus the pooling payoff.

In numerical simulations, as r or σ^2 increases, the payoff of H from an informative equilibrium minus the pooling payoff falls. Intuitively, patience favours signalling and noise favours pooling. Across values of r and σ^2, the payoff difference between an informative equilibrium and pooling can be positive or negative.

For the L type, as r increases, the payoff from an informative equilibrium minus the pooling payoff rises. Since in the signalling region, L expects the receivers' log likelihood ratio (and L's own future payoff) to fall, the more the future payoff matters, the worse off the occurrence of signalling makes L. As σ^2 increases, L's payoff from an informative equilibrium increases—noise is good for L, since the receivers learn about the types more slowly.

So far, only equilibria with effort zero or one have been considered. Other kinds of equilibria also exist: one class is where $e_H^* \in (0,1)$ and $e_L^* = 0$, another class has $e_H^* = 1$ and $e_L^* \in (0,1)$. These interior effort equilibria are more difficult to work with than the equilibria where $e_\theta^* \in \{0,1\}$ and fewer results are available. They are discussed in the online appendix. For some parameter values, there exists a continuum of interior effort equilibria.

As in the one-shot noisy signalling game, a natural question arising in the above Brownian model is whether the results are driven by the linear cost. The online appendix solves a Brownian signalling game with quadratic cost. There is a continuum of equilibrium signalling intervals and on each signalling interval, a continuum of equilibrium effort profiles. The reason why many effort profiles on a given interval constitute equilibria is that the first order conditions are linearly dependent. This is a feature of the quadratic cost and does not generalize to other convex cost functions.

4 Discussion

4.1 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 (an overview of Kaya (2009) is given in the online appendix). In the present paper, the existence of informative equilibria depends on the prior.

Signalling over time has also been studied in continuous time. Daley and Green (2012b), Gryglewicz (2009) and Dilme (2012) use Brownian noise to model imperfect observation of the sender's action, similarly to the Brownian signalling model of this paper. In Dilme (2012), 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 (2012b), 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 (2012b), 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. An example is completing a traditional education—the salary is received only after graduating. In the current paper, the benefit is received concurrently with the payment of the cost, as when a worker takes continuing education courses while being employed, or a firm advertises while selling its product. 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. The signaller could be selling a house or a car, and the buyers can interpret quick agreement as a signal of low quality. In this paper, the sender does not choose whether to trade or not, but exerts a signalling effort, e.g. a firm advertises a product that is already on the shelves of retailers.

One-shot noisy signalling has been studied by Matthews and Mirman (1983), Carlsson and Dasgupta (1997) and Daley and Green (2012a). 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 involving the politician. The motivation for adding noise in Matthews and Mirman (1983) and Daley and Green (2012a) is to better describe real-life signalling situations. Carlsson and Dasgupta (1997) use noise to eliminate unintuitive equilibria in the noiseless model. Another motivation for adding noise in the present paper is to remove equilibria featuring pooling on positive effort, which exist in Spence (1973).

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. Career concerns describe why a manager gets a raise after working hard at his job—the employer cares about both effort and talent. Signalling describes why a manager who completes a (for the sake of the argument, unproductive) MBA gets a raise—the employer only cares about the talent.

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.

The extremal equilibria of the dynamic signalling models of this paper can be interpreted as the low-cost type exerting effort to build a reputation for having low cost. Once signalling stops, belief stops changing, so reputations are in a sense permanent. On the other hand, for beliefs in the signalling region, reputation must constantly be supported at a cost, otherwise it is likely to deteriorate. Therefore in the signalling region, reputation is transitory.

Cripps, Mailath, and Samuelson (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 H may be lower than the prior and belief about L may be higher. In expectation, beliefs move in the direction of the sender's type, but mistakes have positive probability. After signalling stops, an equilibrium of the incomplete information stage game is played forever. In the Poisson models, in extremal equilibria in which the signalling region is infinite, the type is perfectly learned in the limit. This is similar to the result of Cripps, Mailath, and Samuelson (2004).

4.2 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 results) 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 it 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, as in career concerns models, 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 H than from L. 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 in l. 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 l if pooling is expected at high l. If the weight the receivers place on future effort is large enough, the flow benefit of the sender may decrease as l rises towards the pooling region. This does not incentivize H 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, so the same result obtains.

A Proofs omitted from the text

Proof of Lemma 1. $V_\theta(l)$ is bounded above by $\int_0^\infty \exp(-rt)\beta_{\max}dt = \frac{\beta_{\min}}{r} \in \mathbb{R}$ 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 after reaching l for the first time. An optimal strategy gives L the continuation value $V_L(l)$ after 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 $V_L(l)$.

Due to the piecewise continuity of the strategies, if $\exists \hat{l}$ s.t. $e_L(\hat{l}) > 0$, then there exists an open interval $(l_1, l_2) \ni \hat{l}$ s.t. $e_L(l') > 0\ \forall l' \in (l_1, l_2)$. If the set of histories where reaching some $\hat{l}$ satisfying $e_L(\hat{l}) > 0$ has positive probability after l under the optimal strategy, then H can imitate L at a strictly lower cost, getting the same benefit. □

Proof of Proposition 3. 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. 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 $\lambda\left[\frac{\beta_{\max}}{r} - \frac{\beta_{\min}}{r}\right]$. Both types have the unique best response $e = 0$ in this situation. □

Lemma 16. *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, the receivers' posterior probability μ that the sender's type is

H drifts down in expectation for L. The flow benefit is increasing in the posterior: $\hat{\beta}'(\mu) > 0$. The posterior has positive variance, which for a concave $\hat{\beta}$ does not raise the payoff of L. Therefore the continuation payoff from l on in an informative equilibrium starting is below the pooling payoff $\frac{\beta(l)}{r}$ when starting at l. □

Proof of Lemma 4. If $e_L^*(l) > e_H^*(l)$, then L weakly prefers $e = 1$ and H weakly prefers $e = 0$. If $V_H(j(l)) = V_L(j(l)) = k$, then due to $V_H \geq V_L$, the jump in value $k - V_\theta(l)$ after a signal is higher for L. The cost of avoiding jumps is strictly larger for L, because $c_L > c_H$. It cannot be that L prefers to avoid jumps and H prefers to allow them. When pooling occurs after the jump, then $V_H(j(l)) = V_L(j(l)) = \frac{\beta(l)}{r}$.

If $0 < e_L^*(l) < e_H^*(l) = 1$, then $j(l) = -\infty$. It is enough to show V_L is strictly increasing in the region in which $0 < e_L^* < e_H^* = 1$, because then indifference to the jump at one l implies the absence of indifference at any $\hat{l} \neq l$ in the region. Efforts are continuous from the left or right, which rules out the situation where $0 < e_L^*(l) < 1$ at one point, with $e_L^* = 0$ or 1 at neighbouring points.

Since $e_L^* < 1$, V_L is unchanged by switching e_L to 0 throughout the region $(\underline{l}, \bar{l})$ in which $0 < e_L^* < e_H^* = 1$. Then the flow rate of jumps and the flow cost are constant in the region, while the flow benefit is strictly increasing. The only influence that might make V_L decreasing is the continuation value at $\bar{l}$. If $\bar{l} = \infty$, then V_L is strictly increasing in the $0 < e_L^* < e_H^* = 1$ region. With $\bar{l}$ finite and pooling at $\bar{l}$, $V_L(\bar{l}) > V_L(l)$ for l in the $0 < e_L^* < e_H^* = 1$ region, because at $\bar{l}$ the probability of jumping to $l = -\infty$ is zero, the flow benefit is strictly higher and the flow cost is the same as at l. Again, V_L is strictly increasing in the $0 < e_L^* < e_H^* = 1$ region. If at $\bar{l}$, a region where $e_L^* > e_H^*$ starts, then l cannot drift into that region, because l drifts down when $e_L^* > e_H^*$. Therefore l must stay at the boundary $\bar{l}$. As shown previously, it cannot be that $e_L^*(l) = 1$ at some l in equilibrium, so $e_L^*(\bar{l}) < 1$. Then e_L can be switched to 0 at $\bar{l}$ without changing V_L. The flow benefit is higher than in the $0 < e_L^* < e_H^* = 1$ region, the flow cost and jump rate are the same and the jumps go to a value higher than $\frac{\beta_{\min}}{r}$. A higher l means a shorter time until reaching $V_L(\bar{l})$, so again V_L is strictly increasing in the $0 < e_L^* < e_H^* = 1$ region.

If at $\bar{l}$, a region where $e_L^* = 0$, $e_H^* > 0$ starts, then the union of that and the $0 < e_L^* < e_H^* = 1$ region can be taken and the preceding reasoning can be applied at the upper boundary of the union. □

Lemma 17. *In the bad news model, if* $\frac{\beta_{\max}-\beta_{\min}}{r} > \frac{c_H}{\lambda}$*, then there exist equilibria in which for some* $l \in \mathcal{L}(e^*)$*,* $e_H^*(l) \in (0,1)$ *and* $e_L^*(l) = 0$*.*

Proof. Take $l_0 \in \mathbb{R}$ such that $\frac{\beta(l_0)-\beta_{\min}}{r} > \frac{c_H}{\lambda}$. An equilibrium in which $e_H^*(l) \in (0,1)$ and $e_L^*(l) = 0$ on $[l_0, \bar{l}) \subset \mathbb{R}$ will be constructed. Assume $\bar{l} - l_0 = \epsilon > 0$ for ϵ small. The probability of reaching $\bar{l}$ from $[l_0, \bar{l})$ is close to 1, so the payoffs of the types on $[l_0, \bar{l})$ are close to $\frac{\beta(\bar{l})}{r} > \frac{\beta(l_0)}{r}$.

If the market expects $e_H^*(l) = 1$, $e_L^*(l) = 0$ (which implies that l jumps to $j(l) = -\infty$ when a signal occurs at $l \in [l_0, \bar{l})$), then H has the unique best response $e = 1$ at l. If the market expects $e_H^*(l) = e_L^*(l) = 0$ (which implies that $j(l) = l$ when a signal occurs at $l \in [l_0, \bar{l})$), then H has the unique best response $e = 0$ at l. By the continuity of $\beta(\cdot)$ and $j(\cdot)$, there exists $\hat{e} \in (0,1)$ s.t. when the market expects $e_H^*(l) = \hat{e}$, $e_L^*(l) = 0$, then H is indifferent between $e = 1$ and $e = 0$ at l. The same reasoning holds for all points in $[l_0, \bar{l})$, with slightly different $\hat{e}$.

If H is indifferent between $e = 1$ and $e = 0$, then L strictly prefers $e = 0$. □

Proposition 18. *If the following hold*

(a) $\lim_{l \to \underline{l}_1+} V_H(l) - \frac{\beta_{\min}}{r} \geq \frac{c_H}{\lambda}$*,*

(b) $\lim_{l \to \infty} V_L(l) - \frac{\beta_{\min}}{r} \leq \frac{c_L}{\lambda}$*,*

(c) $\frac{\beta(\underline{l})}{r} - \lim_{l' \to \underline{l}_1+} V_H(l') < \frac{c_H}{\lambda}$*,*

(d) $\frac{\beta(\underline{l})}{r} - \lim_{l' \to \underline{l}_1+} V_L(l') > \frac{c_L}{\lambda}$*,*

(e) $\frac{\beta(\underline{l})}{r} - \lim_{l' \to \infty} V_L(l') < \frac{c_L}{\lambda}$*,*

where $-\infty < \underline{l} < l_0 < \underline{l}_1$ *and* V_H, V_L *are given in terms of primitives in (6), then there exists an equilibrium in which*

- $e_L^*(l) > e_H^*(l)$ *if* $l \in (\underline{l}, l_0]$*,*
- $e_L^*(l) = 0$*,* $e_H^*(l) = 1$ *if* $l \in [\underline{l}_1, \infty)$*,*
- $e_L^*(l) = e_H^*(l) = 0$ *if* $l \notin (\underline{l}, l_0] \cup [\underline{l}_1, \infty)$*.*

Proof. The candidate value functions of the types are continuous and strictly increasing on $[\underline{l}_1, \infty)$ according to (6).

For H to optimally choose $e = 1$ in $[\underline{l}_1, \infty)$, it is necessary and sufficient that $V_H(l) - \frac{\beta_{\min}}{r} \geq \frac{c_H}{\lambda}$ for all $l \in [\underline{l}_1, \infty)$. Since V_H is strictly increasing, the inequality holds for all $l \in [\underline{l}_1, \infty)$ iff $\lim_{l \to \underline{l}_1+} V_H(l) - \frac{\beta_{\min}}{r} \geq \frac{c_H}{\lambda}$. A lower bound on $\lim_{l \to \underline{l}_1+} V_H(l)$ is $\frac{\beta(\underline{l}_1) - c_H}{r}$.

For L to optimally choose $e = 0$ in $[\underline{l}_1, \infty)$, it is necessary and sufficient that $V_L(l) - \frac{\beta_{\min}}{r} \leq \frac{c_L}{\lambda}$ for all $l \in [\underline{l}_1, \infty)$. Since V_L is strictly increasing, the inequality holds for all $l \in [\underline{l}_1, \infty)$ iff $\lim_{l \to \infty} V_L(l) - \frac{\beta_{\min}}{r} \leq \frac{c_L}{\lambda}$. The limit is $\lim_{l \to \infty} V_L(l) = \frac{\beta_{\max}}{r + \lambda} + \frac{\lambda \beta_{\min}}{r(r+\lambda)}$.

In $(\underline{l}, l_0]$ in the absence of a signal l drifts down and eventually reaches $\underline{l}$ with positive probability. For any $l, l' \in [\underline{l}, l_0]$ s.t. $l' < l$, the probability of reaching l' from l approaches 1 as $l - l' \to 0$, therefore for any $\epsilon_1 > 0$ there exists $\epsilon_2 > 0$ s.t. $l_0 - \underline{l} < \epsilon_2$ implies $|V_\theta(l) - \frac{\beta(\underline{l})}{r}| < \epsilon_1$ for all $l \in (\underline{l}, l_0]$ and $\theta = 1, 2$.

For all $l \in (\underline{l}, l_0]$, take $e_H^*(l) = 0$ and $e_L^*(l) \in (0, 1)$ such that $j(l) \in [\underline{l}_1, \infty)$, where $j(l) = l + \ln \frac{1 - e_H^*(l)}{1 - e_L^*(l)}$.

For H to optimally choose $e = 0$ at $l \in (\underline{l}, l_0]$, it is sufficient that $V_H(l) - V_H(j(l)) < \frac{c_H}{\lambda}$. Sufficient for this is $V_H(l) - \lim_{l' \to \underline{l}_1+} V_H(l') < \frac{c_H}{\lambda}$. Choose ϵ_2 s.t. $|\frac{\beta(\underline{l})}{r} - V_H(l)| \leq \epsilon_1 \leq \frac{c_H}{\lambda} - V_H(l) + \lim_{l' \to \underline{l}_1+} V_H(l')$. Then sufficient conditions for H to choose $e = 0$ are $l_0 - \underline{l} < \epsilon_2$ and $\frac{\beta(\underline{l})}{r} - \lim_{l' \to \underline{l}_1+} V_H(l') < \frac{c_H}{\lambda}$.

For L to optimally choose $e_L \in (0, 1)$ at $l \in (\underline{l}, l_0]$, it is necessary and sufficient that $V_L(l) - V_L(j(l)) = \frac{c_L}{\lambda}$. By continuity of V_L on $[\underline{l}_1, \infty)$, sufficient for this are $V_L(l) - \lim_{l' \to \underline{l}_1+} V_L(l') > \frac{c_L}{\lambda}$ and $V_L(l) - \lim_{l' \to \infty} V_L(l') < \frac{c_L}{\lambda}$. Then by reasoning similar to the H case in the previous paragraph, sufficient for $e_L \in (0, 1)$ are $l_0 - \underline{l} < \epsilon_2$, $\frac{\beta(\underline{l})}{r} - \lim_{l' \to \underline{l}_1+} V_L(l') > \frac{c_L}{\lambda}$ and $\frac{\beta(\underline{l})}{r} - \lim_{l' \to \infty} V_L(l') < \frac{c_L}{\lambda}$. □

Proof of Proposition 5. If $\frac{\beta_{\max} - \beta_{\min}}{r} > \frac{c_H}{\lambda}$, then $\exists l_0$ for which an informative extremal equilibrium can be constructed. 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 extremal equilibrium. □

Proof of Lemma 7. If $e_L^* = e_H^*$, then the log likelihood ratio stays constant regardless of the occurrence or absence of signals. Then both types optimally choose $e_\theta = 0$ to avoid the effort cost. This rules out $e_L^* = e_H^* > 0$ occurring in equilibrium.

Three steps are needed to rule out $e_L^* > e_H^*$. First, in equilibrium L always has a best response that avoids jumps up. Second, in the region where $e_L^* > e_H^*$, $V_L(l)$ 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 $e_L^* > e_H^*$, then jumps go down. If $e_L^* < e_H^*$, then jumps go up, but $e_L^* < 1$, so $e = 0$ is a best response for L. No jumps occur with $e = 0$.

Step 2. In the region where $e_L^* > e_H^*$, the log likelihood ratio drifts up and jumps down. Taking $e = 0$ 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 l' in the region where $e_L^* > e_H^*$. At $\hat{l}$, L has a best response that avoids jumps. At l', V_L is at least $\frac{\beta(l')}{r}$, which is the payoff to a strategy that avoids jumps. 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 concludes ruling out $e_L^*(l) > e_H^*(l)$.

To rule out $e_L^* = 0 < e_H^* < 1$, consider an interval $(\underline{l}_1, l_0]$ in which $e_L^* = 0 < e_H^* < 1$ is expected. Type H must be indifferent, so switching type H's choice from e_H^* to 0 in the whole $(\underline{l}_1, l_0]$ does not change V_H. If $e = 0$, 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 $l' > l''$. Starting at l' or l'' yields flow cost zero. Starting at l' yields initially 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 $l = \infty$, due to $e_H^* > e_L^* = 0$. So if H is indifferent between $e > 0$ and $e = 0$ 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 $e_L^* = 0 < e_H^* < 1$ 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 one point, with e_H^* either 0 or 1 at neighbouring points cannot occur. □

Lemma 19. *$\forall l_0 \in \overline{\mathbb{R}}\ \forall l \in \mathcal{L}(e^*)$, the following cannot occur:*

(a) $e_L^*(l), e_H^*(l) \in (0,1)\ \forall l \in (\underline{l}, l_0]$, with $\underline{l} > -\infty$ and $e_L^*(\underline{l}) = e_H^*(\underline{l}) = 0$,

(b) $e_L^*(l), e_H^*(l) \in (0,1)\ \forall l \in (\underline{l}, l_0]$, with $e_L^*(j(l)) = e_H^*(j(l)) = 0$ for some $l \in (\underline{l}, l_0]$.

Proof. Lemma 7 ruled out $e_L^*(l) = e_H^*(l) > 0$ and $e_L^*(l) > e_H^*(l)$, so $e_L^*(l), e_H^*(l) \in (0,1)$ implies $e_L^*(l) < e_H^*(l)$ and $j(l) > l$. By Lemma 1, $V_H(j(l)) \geq V_L(j(l))$.

(a) For every $\epsilon > 0$, there exists $\epsilon_2 > 0$ s.t. $|l - \underline{l}| < \epsilon_2$ implies that the probability of reaching $\underline{l}$ from l is $1-\epsilon$. For every $\epsilon_3 > 0$, there exists $\epsilon > 0$ s.t. if the probability of reaching $\underline{l}$ from l is $1-\epsilon$, then $|V_\theta(l) - V_\theta(\underline{l})| < \epsilon_3$. Due to pooling at $\underline{l}$, $V_H(\underline{l}) = V_L(\underline{l}) = \frac{\beta(\underline{l})}{r}$. Combining this with $V_H(j(l)) \geq V_L(j(l))$ and $c_H < c_L$, it is clear that $V_H(j(l)) - V_H(l) = \frac{c_H}{\lambda}$ implies $V_L(j(l)) - V_L(l) < \frac{c_L}{\lambda}$. Both types cannot simultaneously be indifferent to signalling at l close to $\underline{l}$.

(b) If $e_L^*(j(l)) = e_H^*(j(l)) = 0$ and $e_L^*(l), e_H^*(l) \in (0,1)$, then the indifference condition at l is $\frac{\beta(j(l))}{r} - V_\theta(l) = \frac{c_\theta}{\lambda}$ for $\theta = H, L$. Both types are indifferent between their equilibrium effort and zero effort, so payoffs are unchanged by switching the chosen effort to zero, keeping expectations constant. The payoff functions V_θ thus obtained are

$$V_\theta(l) = \exp\left(-\int_{\underline{l}}^{l} \frac{rdz}{\lambda(e_H^*(z) - e_L^*(z))}\right) V_\theta(\underline{l}) + \int_{\underline{l}}^{l} \frac{\beta(x)}{\lambda(e_H^*(x) - e_L^*(x))} \exp\left(-\int_x^l \frac{rdz}{\lambda(e_H^*(z) - e_L^*(z))}\right) dx \tag{10}$$

if $\underline{l} > -\infty$, and

$$V_\theta(l) = \int_{-\infty}^{l} \frac{\beta(x)}{\lambda(e_H^*(x) - e_L^*(x))} \exp\left(-\int_x^l \frac{rdz}{\lambda(e_H^*(z) - e_L^*(z))}\right) dx. \tag{11}$$

if $\underline{l} = -\infty$. The expressions (10) and (11) satisfy the ODE $rV_\theta(l) = \beta(l) - \lambda(e^*_H(l) - e^*_L(l))V'_\theta(l)$, which is derived from the HJB equation using the indifference condition $\lambda[\frac{\beta(j(l))}{r} - V_\theta(l)] - c_\theta = 0$.

Take $l_1, l_2 \in (\underline{l}, l_0]$ with $l_1 < l_2$ and $e^*_\theta(j(l_1)) = e^*_\theta(j(l_2)) = 0$ for $\theta = H, L$. Due to piecewise continuity of the strategies, if there is one point in $(\underline{l}, l_0]$ satisfying $e^*_\theta(j(l)) = 0$ for $\theta = H, L$, then there is a continuum. Since l drifts down in the absence of signals and l_1 is reachable from l_2, $V_\theta(l_1)$ is a continuation value from the perspective of l_2. Based on (10) and (11), $V_\theta(l_2) = \alpha V_\theta(l_1) + k$, with $\alpha \in (0,1)$ and $k \in \mathbb{R}$. Specifically, for both types,

$$\alpha = \exp\left(-\int_{l_1}^{l_2} \frac{rdz}{\lambda(e^*_H(z) - e^*_L(z))}\right),$$
$$k = \int_{l_1}^{l_2} \frac{\beta(x)}{\lambda(e^*_H(x) - e^*_L(x))} \exp\left(-\int_{x}^{l_2} \frac{rdz}{\lambda(e^*_H(z) - e^*_L(z))}\right) dx.$$

Now it can be shown that the indifference conditions for both types for l_1, l_2 contradict each other: subtracting $\lambda[\frac{\beta(j(l_1))}{r} - V_L(l_1)] - c_L = 0$ from $\lambda[\frac{\beta(j(l_1))}{r} - V_H(l_1)] - c_H = 0$, we get $\lambda[-V_H(l_1) + V_L(l_1)] = c_H - c_L$. Subtracting $\lambda[\frac{\beta(j(l_1))}{r} - \alpha V_L(l_1) - k] - c_L = 0$ from $\lambda[\frac{\beta(j(l_1))}{r} - \alpha V_H(l_1) - k] - c_H = 0$, we get $\lambda\alpha[-V_H(l_1) + V_L(l_1)] = c_H - c_L$. This contradicts $\alpha \in (0,1)$, $\lambda > 0$, $c_L > c_H > 0$, which proves (b). □

Lemma 20. *In the good news model, if* $\frac{\beta_{\max} - \beta_{\min}}{r} > \frac{c_L}{\lambda}$*, then there exist equilibria in which for some* $l \in \mathcal{L}(e^*)$*,* $e^*_H(l) = 1$ *and* $e^*_L(l) \in (0,1)$*.*

Proof. Take $l_0 \in \mathbb{R}$ such that $\frac{\beta_{\max} - \beta(l_0)}{r} > \frac{c_L}{\lambda}$. An equilibrium in which $e^*_H(l) = 1$ and $e^*_L(l) \in (0,1)$ on $(\underline{l}, l_0] \subset \mathbb{R}$ will be constructed. Assume $l_0 - \underline{l} = \epsilon > 0$ for ϵ small. The probability of reaching $\underline{l}$ from $(\underline{l}, l_0]$ is close to 1, so the payoffs of the types on $(\underline{l}, l_0]$ are close to $\frac{\beta(\underline{l})}{r} < \frac{\beta(l_0)}{r}$.

If the market expects $e^*_H(l) = 1$, $e^*_L(l) = 0$ (which implies that l jumps to $j(l) = \infty$ when a signal occurs at $l \in (\underline{l}, l_0]$), then L has the unique best response $e = 1$ at l. If the market expects $e^*_H(l) = e^*_L(l) = 1$ (which implies that $j(l) = l$ when a signal occurs at $l \in (\underline{l}, l_0]$), then L has the unique best response $e = 0$ at l. By the continuity of $\beta(\cdot)$ and $j(\cdot)$, there exists $\hat{e} \in (0,1)$ s.t. when the market expects $e^*_L(l) = \hat{e}$, $e^*_H(l) = 1$, then L is indifferent between $e = 1$ and $e = 0$ at l. The same reasoning holds for all points in $[l_0, \bar{l})$, with slightly different $\hat{e}$.

If L is indifferent between $e = 1$ and $e = 0$, then H strictly prefers $e = 1$. □

Proof of Lemma 8. If $l = \infty$ or $l = -\infty$, then l does not respond to signals, so clearly neither type will take positive effort.

The drift in l is finite and the discount rate r is positive, so for any $\epsilon > 0$ $\exists T > 0$ s.t. the flow payoff after time T contributes less than ϵ to total payoff. For any $\epsilon > 0$ and $T > 0$ $\exists \hat{l} \in \mathbb{R}$ s.t. starting at $\hat{l}$, after drifting down at rate λ for length of time T, the log likelihood ratio l reached satisfies $|\frac{\beta(l)-\beta_{\max}}{r}| < \epsilon$. The quantity $V_H(\hat{l})$ is bounded below by the payoff from taking $e = 0$ forever, which makes the rate of jumps zero. The payoff from taking $e = 0$ forever starting at $\hat{l}$ is bounded below by $\frac{\beta(l)}{r} - \epsilon$, where l is reached from $\hat{l}$ after length of time T. Therefore $|V_H(\hat{l}) - \frac{\beta_{\max}}{r}| < 2\epsilon$.

Type H's cost of choosing $e = 1$ over a time interval Δ is $c_H\Delta$ and, starting at $\hat{l}$, the benefit is bounded above by $|V_H(\hat{l}) - \frac{\beta_{\max}}{r}|[1 - \exp(-\lambda\Delta)]$. Thus at $\hat{l}$ there exists $\epsilon > 0$ s.t. the optimal choice of H is $e = 0$. If $e_H^*(l) > 0$ is expected for $l \geq \hat{l}$, then H will deviate to $e = 0$.

If the expectations are $e_L^* \in (0, 1)$, $e_H^* = 1$, then jumps end at some $j(l) < \infty$, which implies a smaller benefit to signalling than in the $e_L^* = 0$, $e_H^* = 1$ case. The previous reasoning still holds, with an even stronger incentive not to signal above $\hat{l}$. □

Proof of Proposition 9. V_H is bounded above by $\frac{\beta_{\max}}{r}$ and below by $\frac{\beta_{\min}}{r}$. If $\frac{\beta_{\max}-\beta_{\min}}{r} \leq \frac{c_H}{\lambda}$, then even jumps from $l = -\infty$ to $l = \infty$ at rate λ do not provide enough benefit to outweigh the cost for H. Thus H will not signal in this case. For any $l \in (-\infty, \infty]$, jumps from $V_H(l)$ to $\frac{\beta_{\max}}{r}$ are smaller than the jumps from $\frac{\beta_{\min}}{r}$ to $\frac{\beta_{\max}}{r}$.

If $\frac{\beta_{\max}-\beta_{\min}}{r} > \frac{c_H}{\lambda}$, then for some l_0 an interval 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 interval equilibrium, so that 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 interval equilibrium. □

Proof of Lemma 11. Due to the boundedness of $\beta(l)$ and e_θ, discounting ensures that V_θ is finite—even without the expectation, the integral in the definition of V_θ is finite for any path of l and any control e_θ.

It is clear that $V_H \geq V_L$ in the signalling region, because H can follow L's strategy at a lower cost than L. Outside the signalling region, $V_H(l) = V_L(l) = \frac{\beta(l)}{r}$.

If there is positive probability of reaching $\hat{l}$ with the optimal $e_L(\hat{l}) > 0$, then H can follow L's strategy at a strictly lower cost.

To prove V_θ is strictly increasing, a standard coupling argument is used. Consider two diffusion processes: the l process with optimal effort starting from l_1 and the l process under zero effort starting from $l_2 > l_1$. Call the former process l^{e^*} and the latter l^0. Define the stopping time $\tau^* = \inf\left\{t > 0 : l_t^0 - l_t^{e^*} = 0\right\}$. The receivers expect the optimal strategy in both cases.

Starting at l_2, the strategy $s =$"play 0 until τ^* and the optimal strategy thereafter" yields a weakly lower payoff than $V_\theta(l_2)$, the payoff to the optimal stationary strategy starting from l_2. This holds even though s is not stationary, because if the receivers expect a stationary strategy, then among the optimal strategies for the sender there is a stationary one. The argument is standard—the competitive receivers always play a static best response, which depends on their belief about the type, but not the sender's strategy, so if at some l, a sender action $\hat{e}$ is optimal at one point in time, then $\hat{e}$ is optimal at that l at another point in time.

Starting at l_2, the strategy s yields a strictly higher payoff than $V_\theta(l_1)$, the payoff to the optimal strategy starting from l_1. This is because the revenue $\beta(l^0)$ is strictly higher than $\beta(l^{e^*})$ before τ^* and the same in expectation after τ^*. The cost of l^0 is zero while the cost of l^{e^*} is positive before τ^*. The costs of the two strategies are the same in expectation after τ^*. Overall, $V_\theta(l_2) > V_\theta(l_1)$ for l_1, l_2 in the signalling region.

If both l_1, l_2 are outside the signalling region, then since β was assumed strictly increasing, the payoffs are ordered $V_\theta(l_2) = \frac{\beta(l_2)}{r} > \frac{\beta(l_1)}{r} = V_\theta(l_1)$. If l_2 is above the signalling region while l_1 is in the signalling region, then the expected benefit is strictly higher from l_2 onwards and the expected cost is the lowest possible from l_2 onwards, so $V_\theta(l_2) > V_\theta(l_1)$. If l_2 is in the signalling region while l_1 is below the signalling region, then $V_\theta(l_2)$ is higher than the payoff to the strategy of taking zero effort forever starting from l_2. The cost of this strategy is the same as the cost of the optimal strategy from l_1 onwards, while the benefit is strictly greater, so again $V_\theta(l_2) > V_\theta(l_1)$. □

Proof of Lemma 12. If $e_H^*(l) = e_L^*(l)$, then the signal is statistically uninformative about the type, so the log likelihood ratio does not respond to the signal. Given this, both types will optimally choose $e_\theta(l) = 0$. Therefore in equilibrium, it cannot be that $e^*_L(l) = e^*_H(l) > 0$ for some l.

A higher expected signal is more costly to both types. If $e^*_L(l) > e^*_H(l)$ is expected, then based on (8), l falls in response to a higher signal. The flow benefit $\beta(l)$ strictly increases in l. By Lemma 11, $V_\theta(l)$ increases in l. If $e^*_L(l) > e^*_H(l)$ is expected, then a higher signal leads to a lower flow benefit, lower continuation value and higher cost, so both types optimally choose $e_\theta(l) = 0$. □

Proof of Proposition 13. L takes no effort in any extremal equilibrium, including pooling, so the flow cost is the same in both cases. The flow benefit comparison is unaffected if $\beta(l)$ is written as $\beta(f^{-1}(f(l)))$ for some strictly increasing smooth f. Use Itō's rule to derive the process $f(l)$:

$$df = \left[\sigma^{-2}(e^*_H - e^*_L)\left(e_L - \frac{1}{2}(e^*_H + e^*_L)\right)\frac{df}{dl} + \frac{1}{2\sigma^2}(e^*_H - e^*_L)^2\frac{d^2 f}{dl^2}\right]dt$$
$$+ \frac{e^*_H - e^*_L}{\sigma}\frac{df}{dl}dB_t.$$

The drift of f is zero iff $\frac{d^2 f}{dl^2} = -2\frac{e_L - \frac{1}{2}(e^*_H + e^*_L)}{e^*_H - e^*_L}\frac{df}{dl}$. Impose the equilibrium condition $e_L = e^*_L$ and recall that in interval equilibria, $e^*_H = 1$ and $e^*_L = 0$. This leads to $\frac{d^2 f}{dl^2} = \frac{df}{dl}$. Using the normalization $f(0) = 1$, $f'(0) = 1$, we get $f(l) = \exp(l)$ and $f^{-1}(z) = \ln(z)$.

If $\beta(\ln(z))$ is concave in z, which has zero drift, then the expectation of $\beta(\ln(z))$ decreases in the variance of z. The variance of z is strictly increasing in the variance of l. In the pooling equilibrium, l is constant, but in informative equilibria, it has positive variance. A similar reasoning establishes that if $\beta(\ln(z))$ is convex in z, then the payoff of L in any informative equilibrium is above the pooling payoff. Analogous reasoning shows H prefers pooling to other extremal effort equilibria iff $\beta(-\ln z)$ is convex in z.

For $l \in \left\{\underline{l}_1, \bar{l}_1\right\}$, it follows from the above that $V_{L1}(l) = \frac{\beta(l)}{r} \geq V_{L2}(l)$.

From any point in $(\underline{l}_1, \bar{l}_1)$, the log likelihood ratio process has positive probability of hitting $\underline{l}_1$ and positive probability of hitting $\bar{l}_1$. The flow cost to L is zero in all extremal equilibria for all l. For the same l, the flow benefit to L is the same in all extremal equilibria. The distribution over paths of l up to hitting $\underline{l}_1$ or $\bar{l}_1$ starting from $l_0 \in (\underline{l}_1, \bar{l}_1)$ is the same in the two equilibria with signalling regions $(\underline{l}_1, \bar{l}_1)$ and $(\underline{l}_2, \bar{l}_2)$, because in both equilibria in the region $(\underline{l}_1, \bar{l}_1)$, H takes action 1 and L takes 0. Therefore the continuation value comparisons $V_{L1}(\underline{l}_1) \geq V_{L2}(\underline{l}_1)$ and $V_{L1}(\bar{l}_1) \geq V_{L2}(\bar{l}_1)$ determine the payoff comparison $V_{L1}(l) \geq V_{L2}(l)$ for any $l \in (\underline{l}_1, \bar{l}_1)$. □

Proof of Lemma 14. For any signalling region $(\underline{l}, \bar{l})$, the solutions of the ODEs are differentiable at least as many times as β on $(\underline{l}, \bar{l})$ and continuous on $[\underline{l}, \bar{l}]$. Since β was assumed twice continuously differentiable, V_L and V_H are as well. Given the signalling region, V_H, V_L are bounded for any path of l and control e_θ. Therefore $V_H(l)$, $V_L(l)$ are integrable in the probability law of the l process that starts from l_0 and is controlled by e_θ, uniformly over Markov controls e_θ. So by Theorem 11.2.2 of Øksendal (2010), V_L, V_H coincide with the value functions V_L, V_H.

Under the previous conditions, Theorem 11.2.3 of Øksendal (2010) shows that the optimal Markov control does as well as the optimal nonanticipating control, so if the receivers expect Markov strategies, then both types of the sender have a Markov best response among their best responses. This does not imply that the payoffs of all non-Markov equilibria can be attained with Markov equilibria, since in a non-Markov equilibrium the receivers expect non-Markov strategies. □

Proof of Proposition 15. The proof first shows that if there exists $l \in \mathbb{R}$ s.t. $\frac{\beta'(l)}{r} > c_H\sigma^2$, then there exists $l_0 \in \mathbb{R}$ s.t. $c_L\sigma^2 > \frac{\beta'(l_0)}{r} > c_H\sigma^2$. An informative extremal equilibrium is constructed with a length-ϵ signalling region $(\underline{l}, \bar{l}) \ni l_0$. For this, the ICs are verified on $(\underline{l}, \bar{l})$ by showing $V_H'(l), V_L'(l)$ are close to $\frac{\beta'(l_0)}{r}$. First, the average slope of $\frac{\beta(l)}{r}$ over $(\underline{l}, \bar{l})$ is close to $\frac{\beta'(l_0)}{r}$. Second, the average slope of V_L, V_H over $(\underline{l}, \bar{l})$ equals the average slope of $\frac{\beta(l)}{r}$. Third, at any l close to l_0, $V_H'(l), V_L'(l)$ are close to the average slope of V_L, V_H over $(\underline{l}, \bar{l})$.

$\beta \in C^2$ and bounded, so $\lim_{l\to-\infty}\beta'(l) = 0 = \lim_{l\to\infty}\beta'(l)$. We know that at some $l \in \mathbb{R}$, $\beta'(l) > c_H r\sigma^2$. By the Mean Value Theorem $\exists\delta > 0$, $\exists l_0 \in \mathbb{R}$ s.t. $c_H r\sigma^2 + 2\delta < \beta'(l_0) < c_L r\sigma^2 - 2\delta$. By $\beta' \in C^1$, the average slope of β over a small interval containing l_0 is close to $\beta'(l_0)$. Formally, there exists $\epsilon > 0$ s.t. for any $\underline{l}, \bar{l}$ satisfying $\underline{l} < l_0 < \bar{l}$ and $\bar{l} - \underline{l} < \epsilon$, we have $c_H\sigma^2 + \frac{\delta}{r} < \frac{\beta(\bar{l})-\beta(\underline{l})}{r(\bar{l}-\underline{l})} < c_L\sigma^2 - \frac{\delta}{r}$.

The candidate value functions V_L, V_H calculated in the main text are infinitely differentiable on the signalling interval $(\underline{l}, \bar{l})$, continuous on $\mathbb{R}$ and satisfy the boundary conditions $V_H(\underline{l}) = V_L(\underline{l}) = \frac{\beta(\underline{l})}{r}$ and $V_H(\bar{l}) = V_L(\bar{l}) = \frac{\beta(\bar{l})}{r}$. By the Mean Value Theorem, there exist $\hat{l}_L, \hat{l}_H \in (\underline{l}, \bar{l})$ satisfying $V_L'(\hat{l}_L) =$

$V_H'(\hat{l}_H) = \frac{\beta(\bar{l})-\beta(\underline{l})}{r(\bar{l}-\underline{l})}$.

By the smoothness of V_H', V_L', if the ICs are satisfied at some $\hat{l}_L, \hat{l}_H \in (\underline{l}, \bar{l})$, then for a small enough $(\underline{l}, \bar{l})$, the ICs are satisfied at all $l \in (\underline{l}, \bar{l})$. Formally, for any $\delta > 0$ there exists $\epsilon_3 > 0$ s.t. if $\bar{l} - \underline{l} < \epsilon_3$, then

$$\max_{l \in (\underline{l}, \bar{l})} \left| V_L'(l) - \frac{\beta(\bar{l}) - \beta(\underline{l})}{r(\bar{l} - \underline{l})} \right| + \left| V_L'(l) - \frac{\beta(\bar{l}) - \beta(\underline{l})}{r(\bar{l} - \underline{l})} \right| < \frac{\delta}{r}.$$

Take $\underline{l}, \bar{l}$ satisfying $\underline{l} < l_0 < \bar{l}$ and $\bar{l} - \underline{l} < \min\{\epsilon_3, \epsilon\}$, then for all $l \in (\underline{l}, \bar{l})$, we have $c_H\sigma^2 < V_H'(l), V_L'(l) < c_L\sigma^2$. The ICs are satisfied, so $(\underline{l}, \bar{l})$ is the signalling region of an interval equilibrium, with $l_0 \in (\underline{l}, \bar{l})$.

If $\nexists l \in \mathbb{R}$ satisfying $c_H r\sigma^2 \leq \beta'(l)$, then there is no signalling interval on which IC_H can be satisfied at every point. The average slope of V_H over any interval is less than $c_H\sigma^2$, so at some l in the interval, $V_H'(l) < c_H\sigma^2$. The maximal benefit to signalling at l occurs when the expectations of the market are $e_H^*(l) = 1$, $e_L^*(l) = 0$. If H cannot be incentivized to signal at these expectations, then no other expectations incentivize H to signal either. Pooling is then the unique equilibrium. □

References

Admati, A. R., and M. Perry (1987): "Strategic delay in bargaining," *The Review of Economic Studies*, 54(3), 345–364.

Carlsson, H., and S. Dasgupta (1997): "Noise-Proof Equilibria in Two-Action Signaling Games," *Journal of Economic Theory*, 77(2), 432 – 460.

Cripps, M. W., G. J. Mailath, and L. Samuelson (2004): "Imperfect monitoring and impermanent reputations," *Econometrica*, 72(2), 407–432.

Daley, B., and B. Green (2012a): "Market Signaling with Grades," Working paper, UC Berkeley.

——— (2012b): "Waiting for News in the Market for Lemons," *Econometrica*, 80(4), 1433–1504.

Dilme, F. (2012): "Dynamic Quality Signaling with Moral Hazard," Discussion paper, 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*, pp. 630–649.

Fudenberg, D., and J. Tirole (1986): "A "signal-jamming" theory of predation," *The RAND Journal of Economics*, 17(3), 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(1), 169–182.

Hörner, J., and N. Vieille (2009): "Public vs. private offers in the market for lemons," *Econometrica*, 77(1), 29–69.

Kaya, A. (2009): "Repeated signaling games," *Games and Economic Behavior*, 66(2), 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(1), 58–82.

Kreps, D. M., and R. Wilson (1982): "Reputation and imperfect information," *Journal of economic theory*, 27(2), 253–279.

Liptser, R. S., and A. N. Shiryaev (1977): *Statistics of random processes I: General theory.* Springer-Verlag, New York.

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*, pp. 981–996.

Milgrom, P., and J. Roberts (1982): "Predation, reputation, and entry deterrence," *Journal of economic theory*, 27(2), 280–312.

Nöldeke, G., and E. van Damme (1990): "Signalling in a dynamic labour market," *The Review of Economic Studies*, 57(1), 1–23.

Øksendal, B. (2010): *Stochastic differential equations. An introduction with applications.* Springer-Verlag, Heidelberg, Germany.

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(3), pp. 355–374.

Swinkels, J. M. (1999): "Education signalling with preemptive offers," *The Review of Economic Studies*, 66(4), 949–970.

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

Yushkevich, A. A. (1980): "On reducing a jump controllable Markov model to a model with discrete time," *Theory of Probability & Its Applications*, 25(1), 58–69.

——— (1988): "On the two-armed bandit problem with continuous time parameter and discounted rewards," *Stochastics*, 23(3), 299–310.