Reputation with Strategic Information Disclosure

Harry Di PEI*

Aug 10th, 2016

Abstract

I study the dynamics of an agent's reputation for competence when the labour market's information about his performance is disclosed by an intermediary who cannot commit. I show that this game admits a unique Markov Perfect Equilibrium (MPE). When the agent is patient, his effort is inverse U-shaped, while the rate of information disclosure is decreasing over time. I illustrate the inefficiencies of the unique MPE by comparing it with the equilibrium in the benchmark scenario where the market automatically observes all breakthroughs. I characterize a tractable subclass of non-Markov Equilibria and explain why allowing players to coordinate on payoff-irrelevant events can improve efficiency on top of the unique MPE and the exogenous information benchmark. When the intermediary can commit, her optimal Markov disclosure policy has a deadline, after which no breakthrough will be disclosed. However, deadlines are not incentive compatible in the game without commitment, illustrating a time inconsistency problem faced by the intermediary. My model can be applied to professional service industries, such as law and consulting. My results provide an explanation to the observed wage and promotion patterns in Baker, Gibbs and Holmström (1994).

Keywords: reputation, information disclosure, Poisson good news

JEL Codes: D82, D83, M52

1 Introduction

Reputation concern is an important driver of incentives in many professions, ranging from consultants, lawyers, judges and fund managers, to scientists, scholars and professional athletes. In particular, the incentive to *establish a name* is more substantial in the early stages of one's career: junior people work

*Department of Economics, Massachusetts Institute of Technology. Email: harrydp@mit.edu. I am indebted to Daron Acemoglu, Drew Fudenberg and Juuso Toikka for guidance and support. For useful comments, I thank Abhijit Banerjee, Sydnee Caldwell, Gonzalo Cisternas, Mira Frick, Robert Gibbons, Sinem Hidir, Bengt Holmström, Johannes Hörner, Irina Kholodenko, Yucheng Liang, Jean Tirole, Muhamet Yildiz, seminar participants at Harvard, MIT, Stanford, Toulouse, Warwick, and in particular, Alessandro Bonatti, Daria Khromenkova and Alexander Wolitzky. All errors are mine.

hard at entry-level jobs, hoping their professions will recognize their talents. However, the labour market rarely receives information about the performance of these lower-level people. Moreover, this information is mostly revealed by their current employers or direct supervisors.[1] This raises concerns that the latter might have incentives to manipulate the market's expectations by releasing information *strategically*.

Motivated by these phenomena, I analyze a reputation building model with the innovation that the market's information about an agent's performance is released strategically by an informational '*intermediary*', who also has a private interest in the game. Different from other dynamic information disclosure models, the intermediary *cannot commit* to disclosure policies.[2] This generates interesting dynamic interactions between the intermediary's incentive to release information and the agent's incentive to build up his reputation.

Related circumstances abound, especially in professional service industries, where team work among junior people makes it hard for the labour market to infer each individual's contribution. As a result, employers (or supervisors) enjoy substantive discretion in revealing information about their subordinates. Consider the example of a law firm, where junior associates work hard on cases but rarely have the opportunity to present their results in court. Their talents are recognized only after they have been given such chances or being promoted. In European soccer clubs (or in other professional sports), youth team players want to impress first team managers in order to gain opportunities to play in higher level matches (such as the Premier League). But the person who monitors their training and knows their abilities is a youth team coach. Similar stories happen between junior consultants and consulting firms, research assistants and professors, judicial clerks and judges, etc.

I examine a continuous time game between an agent (he), an intermediary (she) and a competitive labour market. The agent's wage is determined by the market's willingness to pay for his service, and the latter depends on his talent, which is his private information. Signals about his performance take the form of *conclusive good news* (or '*breakthroughs*') and arrive according to a Poisson process, with arrival rate increasing in the agent's talent as well as his unobserved effort.[3] Moreover, talent and effort

[1] Asymmetric learning between the current employers and the labour market about a worker's ability is a well-known fact documented in the labour economics literature, including the theoretical works of Waldman (1984), Milgrom and Oster (1987) etc. In a recent paper, Kahn (2013) uses National Longitudinal Survey of Youth data and finds empirical evidence supporting asymmetric learning.

[2] Examples where the intermediary can commit include Ekmekci (2011), Ely (2015), Halac et.al.(2015), Kremer et.al.(2015), Che and Hörner (2015), Hörner and Lambert (2015), etc.

[3] As argued in Bonatti and Hörner (2015), learning via infrequently arrived good news is a distinctive feature of '*creative industries*', in which '*creativity and originality are essential for success*'. This applies to professional service industries, namely, law and consulting, R&D, academia, professional sports, etc. in which a reputation is established at several defining moments of one's career. This includes, for example, bringing in a new client to the firm, making a breakthrough innovation, publishing a paper, scoring a hat-trick in an important match, etc.

are complements (Dewatripont et.al.[1999]).

The novelty of my model is that the market observes a breakthrough only after the intermediary discloses it. Motivated by the aforementioned applications, I assume that the intermediary benefits from the agent's effort as well as from establishing him in front of the public.[4] Since one breakthrough is sufficient to convince the market about the agent's competence, no interesting strategic interaction takes place once a breakthrough is disclosed, after which both the intermediary and the agent receive a fixed continuation value.

In section 3, I characterize the *unique* Markov Perfect Equilibrium (MPE) of this game. I find that when players are patient, the agent's effort is inverse U-shaped and the disclosure rate is decreasing over time. Intuitively, since the intermediary cannot commit, her incentive to disclose information only depends on the comparison between the future revenue she can milk from the agent and the lump sum payoff from disclosure. Hence, she has more incentive to withhold information when the agent's future effort is higher. On the other hand, since the market learns via infrequently arrived good news, the agent's incentive to exert effort decreases with his current continuation value (which is increasing in his reputation) and increases with the current disclosure rate. As a result, the intermediary has an incentive to suppress information only when the market becomes pessimistic and the rate of disclosure is decreasing over time. In response to this, the agent's effort eventually decreases. My result predicts that conditional on staying at the entry level job, both the agent's real wage as well as his chances of being promoted are inverse U-shaped, which echoes the empirical findings in Baker, Gibbs and Holmström (1994a,b).[5]

Comparing this with the unique equilibrium in the benchmark scenario in which the market observes all breakthroughs automatically (or *exogenous information benchmark*), I find that in both cases, the agent's effort is too low at the beginning, relative to social first best, exhibiting a *procrastination inef-*

[4]The intermediary benefits from establishing her agent is a distinct and indispensable feature of my model, which is relevant in the aforementioned applications. In professional sports, soccer clubs receive transfer fees by selling their players to bigger clubs. In professional service industries, namely law and consulting, an established former employee is more likely to work in-house for a potential client, rather than working for a competing firm in the same industry. Once an alumni becomes an in-house attorney or counsel for a client, this also creates a network benefit for his former law or consulting firm.

For example, a former Boston Consulting Group (BCG) employee, who now runs Red Hat, a big software company, said that "*I wouldn't say I am blindly loyal, but I do use BCG more than any other firm.*"(excerpted from an Economist article in 2014) An SRZ (a large law firm in NYC) alumni, who is now the General Counsel at Sterling Stamos Capital Management, said in an interview that "*so while you need to have great knowledge of the law, you also need great resources. This is where SRZ comes in...*" (excerpted from SRZ alumni website) Network benefits can also take various other forms. For example, an article about professional service firms in *The Economist* suggests that "*... former employees are increasingly treated as assets, not turncoats... such firms are trying to stay in touch with departed workers, hoping to turn them into brand ambassadors, recruiters and salespeople*".

[5]Baker, Gibbs and Holmström (1994a,b) use 20 years of personnel data of management employees from a large US firm and find these two empirical facts. First, the real wage of a worker conditional on no promotion is first increasing and then decreasing over time (Figure IV, page 951). Second, the promotion rate from level 2 (entry level for management employees) to level 3 (intermediate level) is also an inverse U-shaped function of time stayed in level 2 (Table IV, page 902).

ficiency. Interestingly, the agent's continuation value is not monotone in disclosure rate. In particular, when effort cost is low, despite the intermediary will withhold information when the market becomes pessimistic, the agent's continuation value at optimistic beliefs can still be higher under endogenous information relative to the exogenous information benchmark, i.e. withholding information can encourage procrastination.

Intuitively, this is because low disclosure rate has two effects on the agent's payoff. A direct effect which makes it hard for the agent to establish himself. An indirect effect which *slows down market learning*. When effort cost is low, the intermediary's equilibrium disclosure rate is also low, and the agent can maintain a good reputation for a long time, although no breakthrough has been disclosed. This is because the market attributes the lack of news to the low disclosure rate, instead of the agent's incompetence. As a result, the agent can sustain a high flow payoff, which discourages him to exert effort.

In section 4, I explore the possibility of mitigating the procrastination inefficiency by allowing players' strategies to depend on payoff irrelevant state variables, for example, whether breakthroughs have been concealed in the past or not. In particular, I characterize a tractable subclass of (non-Markov) Perfect Bayesian Equilibrium, which I call Semi-Markov Equilibrium (SME), in which players' strategies are only required to be Markov on the equilibrium path. This solution concept minimally departs from MPE and the equilibrium strategies have intuitive interpretations: every SME is characterized by a *cutoff belief*, such that the intermediary withholds information only when the market's belief falls below this cutoff. This cutoff can take any value within a compact interval, with the unique MPE being the SME with the highest cutoff and the equilibrium under exogenous information being the SME with the lowest cutoff.

Since the time at which the intermediary starts to withhold information matters for the agent's incentive, SME improves efficiency by offering flexibility in choosing this cutoff belief. For example, when effort cost is low, in order to prevent the agent from receiving high flow payoffs before establishing himself, withholding information should only happen at sufficiently pessimistic beliefs.

In section 5, I characterize the optimal Markov policy when the intermediary can commit. When the market's prior belief is not too pessimistic,[6] there exists a '*deadline*' in the optimal policy, after which no information is disclosed and the agent has no incentive to exert effort. This is in sharp contrast to the no commitment case, in which the incentive to exert effort never vanishes before a breakthrough is disclosed. This deadline minimizes the agent's continuation value upon reaching it, which facilitates incentive provision early on in the game. However, due to the intermediary's sequential rationality constraint, deadlines

[6]Notice that the optimal Markov policy depends on the market's prior belief.

are never implementable in any Perfect Bayesian Equilibrium in absence of commitment.

In section 6, I examine the robustness of my result along several directions, including cases in which the intermediary can disclose past breakthroughs, the market can also learn from public signals in addition to the intermediary's private signals, the intermediary is the agent's direct supervisor instead of his current employer, etc. I also discuss how to enrich my model in order to incorporate more realistic features.

Related Literature: This paper contributes to a burgeoning literature on the dynamic provision of incentives via information disclosure in principal-agent relationships. The majority of these papers focus on how to motivate an agent by providing him informational feedback. Some prominent examples include Campbell et.al.(2014), Ely (2015), Kremer et.al.(2015), Che and Hörner (2015), Halac et.al.(2015). Comparing with these papers, I ask a novel question, that is, how to motivate an agent by releasing information to a *third party* (namely, the market).

Two papers that share a similar motivation with mine are Ekmekci (2011) and Hörner and Lambert (2015). The former constructs a rating system that sustains reputation building incentives. The latter studies the Holmström (1999) career concern model and examines the optimal design of public information structure to maximize the agent's effort in stationary equilibria.

The main differences are as follows. First, in terms of the modeling choice, the informational intermediary can commit to dynamic disclosure policies in both of these papers, which assumes away the strategic issues in information disclosure. In contrast, I focus on the complementary case in which the intermediary *cannot commit* and is a strategic long run player. While the commitment benchmark fits better into applications such as online platforms and credit rating agencies, the non-commitment case is more coherent to disclosure problems in firms and organizations, where employers enjoy substantive discretion in releasing information about their employees. Second, in terms of the main result, I focus on the *dynamics* of effort and disclosure rate, instead of examining the sustainability of reputation building incentives or the maximal *stationary* effort level.

My model builds on continuous time reputation models with Poisson good news, such as Faingold and Sannikov (2011).[7] The main difference is that there are two long run players, giving rise to a multiplicity of equilibria. As a result, more restrictive solution concepts (MPE and SME) are required to make sharp predictions. The intertemporal substitutability of effort is also reported in the strategic experimentation

[7] Other reputation models with Poisson good news include Board and Meyer-ter-Vehn (2013), Halac and Prat (2015), etc. After removing the intermediary, my model can be viewed as a continuous time analogue of Mailath and Samuelson (2001), in which the competent agent exerts effort to distinguish himself from the inept type.

literature as well as its applications in venture capital financing (for example, Bergemann and Hege [1998, 2005], Hörner and Samuelson [2013]). Common in these papers, if an agent shirks today, he retains the option value of succeeding tomorrow, which causes inefficient delays. These papers offer solutions when formal contracts are available: by making the agent's share of surplus a decreasing function of the elapsed time. My paper proposes a complementary solution when output is not contractible: by decreasing the publicity of the agent's performance after the market's belief falls below an endogenously chosen cutoff.[8]

2 The Baseline Model

I introduce a baseline model in this section, which highlights the mechanisms at work and will be the main focus of this paper. I will discuss the robustness of my results as well as how to enrich this model to incorporate more realistic features in section 6.

Players & Actions: There is an agent (he, junior worker), an intermediary (she, current employer of the worker) and a competitive labour market (or '*market*'). Time $t \in [0, +\infty)$ is continuous. The agent's type is denoted by $\theta \in \{0, 1\}$, which is constant over time and is either high ($\theta = 1$) or low ($\theta = 0$). At every time t, he chooses an effort level $\tilde{a}_t \in [0, 1 - \phi]$, where $\phi \in (0, 1)$, and produces outputs called '*breakthroughs*', which are generated according to a Poisson process with arrival rate $\mu\theta(\tilde{a}_t + \phi)$, where $\mu > 0$ is a parameter. As we can see, effort and talent are complements (Dewatripont et.al.[1999]).

Whenever a breakthrough arrives, the intermediary decides at that instant between disclosing it publicly and withholding it.[9] Importantly, the intermediary *cannot* commit to dynamic disclosure policies and *cannot* disclose when breakthroughs do not exist.[10] Let $\chi_t \in [0, 1]$ be the probability of disclosing a breakthrough conditional on its arrival at time t. Throughout the paper, I will be focusing on the '*informational intermediary*' role of the employer while abstracting away from the others.

Information Structure & History: The agent's effort and his type are only known by himself, i.e unbeknownst to the market and the intermediary. Let $\pi_0 \in (0, 1)$ be the probability that their prior

[8] Another difference between these papers and mine is that the agent does not know his type in strategic experimentation models, while in my reputation building model, he knows his type. I will discuss the '*career concern*' case in section 6.

[9] As I will explain later, the equilibrium in my baseline model remains robust when we allow for disclosing past breakthroughs (or '*delayed disclosure*'). Moreover, in other equilibria that arise under delayed disclosure, several important qualitative features of the equilibrium in the baseline model remain robust. I will discuss this in section 6.

[10] Scenarios in which the intermediary cannot announce '*forged breakthroughs*' include: when information is verifiable, or when she faces significant penalties (for example, in terms of reputation costs) for disclosing false information, etc. See Milgrom (2008) for more discussions on the applications of disclosure models.

belief attaches to $\theta = 1$. Let π_t be *the market's* posterior belief at time t, which I will refer to as '*the agent's reputation*'. Whenever a breakthrough arrives, it is automatically observed by the agent and the intermediary. The novelty of my model is that the market can observe a breakthrough *if and only if* the intermediary discloses it. Since the low type can never produce any breakthroughs, the market knows $\theta = 1$ after a disclosure. I say that the agent '*establishes himself*' when π_t reaches 1.

A public history, $h^t \in [0,t]$, consists of a sequence of past disclosure dates $0 \leq t_1 < ... < t_n \leq t$, with $\{\varnothing\}$ the history that no breakthrough has been disclosed. The intermediary's private history, $h^t_m \equiv (h^t, \overline{h}^t_m)$, consists of the public history, as well as $\overline{h}^t_m \subset [0,t]$, which is a sequence of breakthrough arrival times. Since breakthroughs cannot be forged, $h^t \subset \overline{h}^t_m$. The agent's private history, $h^t_a \equiv (h^t_m, a^t, \theta)$, consists of the intermediary's private history, his past effort choices, $a^t \equiv \{a_{t'}\}_{t' \in [0,t]}$, as well as his type. Let h^{t-}, h^{t-}_m and h^{t-}_a be the (public and private) histories up to, but not including time t, and let H^{t-}, H^{t-}_m and H^{t-}_a be the set of histories up to time t. Let $H \equiv \bigcup_{t=0}^{\infty} H^t$, $H_m \equiv \bigcup_{t=0}^{\infty} H^t_m$ and $H_a \equiv \bigcup_{t=0}^{\infty} H^t_a$ be the entire set of histories. I use '$\succ$' to denote the successor relationship between two histories.

Payoffs: Both the agent and the intermediary are risk neutral and discount their future payoffs at rate $r > 0$. For a player receiving flow payoffs $\{U_t\}_{t \geq 0}$, his or her (normalized) *continuation value* at t is:

$$r \int_t^{\infty} e^{-r(s-t)} U_s ds.$$

As in other reputation or career concern models, for example, Holmström (1999), Mailath and Samuelson (2001), I assume that the agent's wage, w_t, is determined by the labour market's expected willingness to pay for his output produced at time t.

The agent's flow payoff is $w_t - c\tilde{a}_t$, where $c \in (0,1)$ is his marginal cost of effort. Since the low type can never produce any breakthroughs, his effort is always 0. Hereafter, I will focus exclusively on the high type. The market's willingness to pay for each breakthrough is $\frac{1}{\mu}$, which normalizes the high type's marginal product of effort to 1 and implies that it is always socially efficient for him to exert effort. Let $a_t \equiv \tilde{a}_t + \phi \in [\phi, 1]$, we have $w_t = \pi_t a_t$.[11] Hereafter, I abuse terminology and refer a_t as the agent's effort.

Next, I specify the intermediary's payoff. According to the interpretation that she is the agent's current employer in professional service industries, she receives the agent's output and pays his wage when $\pi_t < 1$, i.e. her flow payoff is $\theta a_t - w_t$ before the agent establishes himself.[12] After $\pi_t = 1$, the

[11] As in other reputation or career concern models, w_t depends on the market's expected a_t, instead of the true a_t. I will distinguish between these two in the '*Strategies*' paragraph, in which I formally introduce the notation for '*believed effort*'.

[12] Her valuation for the agent's output (or breakthrough) is the same as the market's.

intermediary's flow payoff is constantly b. In the law and consulting firm application, the agent leaves the entry level job offered by the intermediary after establishing himself, and b is her network benefit for having one of her former employees working in-house for a potential client. In what follows, I will focus on the case in which $b \in (\phi, 1)$.[13] I will discuss other values of b as well as alternative specifications of the intermediary's payoff in section 6.

Strategies: The agent chooses an *effort plan*, $\mathbf{a} \equiv \{a(h_a^t)\}_{h_a^t \in H_a}$. The intermediary chooses a *disclosure plan* $\boldsymbol{\chi} \equiv \{\chi(h_m^t)\}_{h_m^t \in H_m}$, with $\chi(h_m^{t-})$ being the probability of disclosing information at time t conditional on $t \in h_m^t$. Truncating the effort plan and the disclosure plan at t, we get $\{a(h_a^{t'}), \chi(h_m^{t'})\}_{t' \leq t}$. This together with the event $\{\theta = 1\}$ induce a probability measure over $H^t \times H_m^t \times H_a^t$, which I denote by $\mathcal{P}_t^{\mathbf{a},\boldsymbol{\chi}}$. The projection of $\mathcal{P}_t^{\mathbf{a},\boldsymbol{\chi}}$ on H^t induces a probability measure over the public histories, with $\mathbb{E}_t^{\mathbf{a},\boldsymbol{\chi}}[\cdot]$ the expectation taken under this measure.

The market's believed effort plan and disclosure plan are $\hat{\mathbf{a}} \equiv \{\hat{a}(h_a^t)\}_{h_a^t \in H_a}$ and $\hat{\boldsymbol{\chi}} \equiv \{\hat{\chi}(h_m^t)\}_{h_m^t \in H_m}$ respectively. These together with π_0 govern the joint distribution over θ and the public histories. Let $\mathcal{P}_t^{\hat{\mathbf{a}},\hat{\boldsymbol{\chi}},\pi_0}$ be the probability measure over $H^t \times H_m^t \times H_a^t$ induced by $\hat{\mathbf{a}}$, $\hat{\boldsymbol{\chi}}$ and π_0. Let $\mathcal{P}_t^{\hat{\mathbf{a}},\hat{\boldsymbol{\chi}},\pi_0}[h^t] \in \Delta(H_m^t \times H_a^t)$ be the projection of $\mathcal{P}_t^{\hat{\mathbf{a}},\hat{\boldsymbol{\chi}},\pi_0}$ on $H_m^t \times H_a^t$ conditional on the public history being h^t, which is the market's '*conditional belief*' over the private histories. Let $\mathbb{E}_t^{\hat{\mathbf{a}},\hat{\boldsymbol{\chi}},\pi_0}[\cdot|h^t]$ be the expectation under $\mathcal{P}_t^{\hat{\mathbf{a}},\hat{\boldsymbol{\chi}},\pi_0}[h^t]$. Let $\boldsymbol{\pi} : H \to \Delta(\Theta)$ be a '*market belief system*', with $\pi(h^t)$ being the probability the market attaches to $\theta = 1$ after observing public history h^t.

Policies: A '*policy*', $(\mathbf{a}, \boldsymbol{\chi})$, consists of an effort plan and a disclosure plan. I introduce two classes of policies: Markov and Semi-Markov Policies.

Definition 1. *A policy is Markov if for every (h_a^{t-}, h_m^{t-}) and $(\hat{h}_a^{t-}, \hat{h}_m^{t-})$, if $\pi(h^{t-}) = \pi(\hat{h}^{t-})$, then*

$$\Big(a(h_a^{t-}), \chi(h_m^{t-})\Big) = \Big(a(\hat{h}_a^{t-}), \chi(\hat{h}_m^{t-})\Big).$$

When analyzing Markov policies, I use $(a(\pi_t), \chi(\pi_t))$ or (a_t, χ_t) instead of $(a(h_a^{t-}), \chi(h_m^{t-}))$ for notation simplicity, where π_t is the market's belief at t when $h^{t-} = \{\varnothing\}$. Furthermore, $\mathbb{E}_t^{\mathbf{a},\boldsymbol{\chi}}[\cdot]$ and $\mathbb{E}_t^{\hat{\mathbf{a}},\hat{\boldsymbol{\chi}},\pi_0}[\cdot|h^{t-} = \{\varnothing\}]$ are induced by Poisson processes with instantaneous arrival rates $\mu a_t \chi_t$ and $\mu \pi_t a_t \chi_t$.

Next, I define Semi-Markov Policies.

[13] When $b \in (\phi, 1)$, conditional on knowing $\theta = 1$, the intermediary has a strict incentive to disclose information if the agent's future effort is constantly ϕ (at its minimum) and has a strict incentive to withhold information if the agent's future effort is constantly 1.

Definition 2. *A policy is Semi-Markov if for every* (h_a^{t-}, h_m^{t-}) *and* $(\hat{h}_a^{t-}, \hat{h}_m^{t-})$*, if* $\pi(h^{t-}) = \pi(\hat{h}^{t-})$ *and* $\mathcal{P}_t^{\mathbf{a},\boldsymbol{\chi}}(h_a^{t-}, h_m^{t-}) > 0, \mathcal{P}_t^{\mathbf{a},\boldsymbol{\chi}}(\hat{h}_a^{t-}, \hat{h}_m^{t-}) > 0$*, then*

$$\Big(a(h_a^{t-}), \chi(h_m^{t-})\Big) = \Big(a(\hat{h}_a^{t-}), \chi(\hat{h}_m^{t-})\Big).$$

In a nutshell, a Semi-Markov policy only requires that $\mathbf{a}$ and $\boldsymbol{\chi}$ to be Markov *on the equilibrium path*, while allowing for non-Markov plans off-path. By definition, the set of Semi-Markov policies contains the set of Markov policies.

Solution Concepts: A *Perfect Bayesian Equilibrium* (or PBE) consists of an equilibrium policy $\{\mathbf{a}, \boldsymbol{\chi}\}$, the market's conditional belief over private histories, $\mathcal{P}_t^{\hat{\mathbf{a}},\hat{\boldsymbol{\chi}},\pi_0}[h^t]$ (induced by the believed policy $\{\hat{\mathbf{a}}, \hat{\boldsymbol{\chi}}\}$) and a wage process $\mathbf{w} : H \to \mathbb{R}_+$, such that:

1. $\mathbf{w}$ is consistent with the market's belief, i.e. $w_t \equiv \mathbf{w}(h^{t-}) = \mathbb{E}_t^{\hat{\mathbf{a}},\hat{\boldsymbol{\chi}},\pi_0}[\theta a(h_a^{t-})|h^{t-}]$.
2. $a(h_a^{t-})$ is optimal for the agent for every h_a^{t-} given $\boldsymbol{\chi}$, $\mathcal{P}_t^{\hat{\mathbf{a}},\hat{\boldsymbol{\chi}},\pi_0}[h^t]$ and $\mathbf{w}$.
3. $\chi(h_m^{t-})$ is optimal for the intermediary for every h_m^{t-} given $\mathbf{a}$, $\mathcal{P}_t^{\hat{\mathbf{a}},\hat{\boldsymbol{\chi}},\pi_0}[h^t]$ and $\mathbf{w}$.
4. For every $h^t \in H$, $\mathcal{P}_t^{\hat{\mathbf{a}},\hat{\boldsymbol{\chi}},\pi_0}[h^t]$ is derived from $\mathcal{P}_t^{\hat{\mathbf{a}},\hat{\boldsymbol{\chi}},\pi_0}$ according to Bayes Rule.[14]
5. The market's belief is correct, i.e. $\{\mathbf{a}, \boldsymbol{\chi}\} = \{\hat{\mathbf{a}}, \hat{\boldsymbol{\chi}}\}$.

Optimality requirements in 2 and 3 imply that at every private history, the agent and the intermediary choose a_t and χ_t respectively to maximize their own expected continuation values, which are given by:

$$\mathbb{E}_t^{\mathbf{a},\boldsymbol{\chi}}\Big[r\int_t^{\infty} e^{-r(s-t)}\Big(w_s - c(a_s - \phi)\Big)ds\Big|h_a^{t-}\Big]$$

and

$$\mathbb{E}_t^{\mathbf{a},\boldsymbol{\chi}}\Big[r\int_t^{\infty} e^{-r(s-t)}U_{m,s}ds\Big|h_m^{t-}\Big],$$

where $U_{m,s} = \theta a_s - w_s$ if $h^s = \{\varnothing\}$ and $U_{m,s} = b$ otherwise. The intermediary's payoff is evaluated conditional on $\theta = 1$ since she has no decision to make until she knows that $\theta = 1$.[15]

[14] Bayes Rule always applies since all public histories occur with strictly positive probability on the equilibrium path.

[15] Formally, since a_t is unobservable to the intermediary, we also need to specify the intermediary's believed effort plan $\overline{a}(h_a^{t-})$, and evaluate her payoff under the probability measure induced by $\{\overline{\mathbf{a}}, \boldsymbol{\chi}\}$ and $\theta = 1$. However, since first, $\mathbf{a} = \overline{\mathbf{a}}$ in equilibrium; and second, the intermediary only moves when her private belief is 1, after which her belief updating process is trivial, so omitting $\overline{\mathbf{a}}$ and evaluating the intermediary's payoff using $\mathbf{a}$ is without loss. I hope this '*inconsistency*' will not cause confusion. However, I need to specify the market's believed effort since it matters for its belief updating process.

Two PBEs are '*outcome equivalent*' if they induce the same joint distribution over $\{\pi_t, w_t, a_t\}_{t\in[0,+\infty)}$. I introduce two refinements of PBE, which will be the solution concepts examined in this paper:

Definition 3. *A Markov Perfect Equilibrium (or MPE) is a PBE in which the equilibrium policy is Markov.*[16] *A Semi-Markov Equilibrium (or SME) is a PBE in which the equilibrium policy is Semi-Markov.*

Let $(\hat{\mathbf{a}}, \hat{\boldsymbol{\chi}}) \equiv \{\hat{a}(\pi_t), \hat{\chi}(\pi_t)\}_{\pi_t\in(0,1)}$ be the believed policy in an MPE, or the believed on-path policy in an SME. I make the technical restriction that both $\hat{a}(\pi_t)$ and $\hat{\chi}(\pi_t)$ are left-continuous functions of π_t. This implies that $a(\pi_t)$ and $\chi(\pi_t)$ are also left-continuous since the market's belief is correct in equilibrium. Let $\mathbf{A}$ and $\mathbf{X}$ be the set of left-continuous Markov effort plans and disclosure plans.

3 Markov Perfect Equilibrium

In this section, I characterize the unique MPE of this game. I show that effort is inverse U-shaped and the disclosure rate is decreasing over time given that no breakthrough has been disclosed in the past. I highlight the agent's and the intermediary's incentives through their best response correspondences in subsection 3.1. I characterize the unique equilibrium when the market can automatically observe all breakthroughs (or the '*exogenous information benchmark*') in subsection 3.2. I state the main characterization result (Proposition 2) and discuss its implications in subsection 3.3, which is also compared with the exogenous information benchmark in subsection 3.5. I prove Proposition 2 in subsection 3.4.

3.1 Preliminaries

Belief Updating: The first step is to specify the evolution of π_t. Similar to Poisson bandit models, for example Keller, Rady and Cripps (2005), if a breakthrough is disclosed, π_t jumps to 1. Otherwise, the evolution of π_t is characterized by the following ordinary differential equation (ODE):[17]

$$\dot{\pi}_t = -\pi_t(1-\pi_t)Y_t, \tag{3.1}$$

[16] Although there has been no agreed upon definition of MPE when actions are unobservable, my definition is in the spirit of Maskin and Tirole (2001) in which players' strategies are only conditioned on the coarsest information partition such that if all other players use measurable strategies, each player's decision-making problem is also measurable. This is partly because the agent knows his own type and the intermediary only acts after knowing the agent's type, which shuts down the channel for private learning.

[17] Admissibility requires that both $a(\pi_t)$ and $\chi(\pi_t)$ are left continuous, so is $\mu a(\pi_t)\chi(\pi_t)$. In a good news model, no news is bad news, i.e. $\dot{\pi}_t \leq 0$, so according to Klein and Rady (2011), there exists a solution to ODE (3.1) for any given π_0. If there are multiple solutions to this initial value problem, I select the one that is consistent with discrete time approximation, which is shown to be unique.

where $Y_t \equiv \mu\chi_t a_t$ is the arrival rate of publicly disclosed breakthroughs conditional on $\theta = 1$.

The Agent's Incentives: Let $V_a(\pi_t)$ be the agent's (equilibrium) continuation value when his reputation is π_t, which can be decomposed into a weighted average of his flow payoff in the time interval $(t, t+dt]$ and his continuation value at $t+dt$:[18]

$$V_a(\pi_t) = r\Big(\pi_t a_t - c(a_t - \phi)\Big)dt + (1 - rdt)\Big(\underbrace{Y_t dt}_{\text{prob. of disclosure}} V_a(1) + \underbrace{(1 - Y_t dt)}_{\text{prob. of no disclosure}} V_a(\pi_{t+dt})\Big). \tag{3.2}$$

Expanding this equation and ignoring higher order terms, we get:

$$V_a(\pi_t) = \Big(\pi_t a_t - c(a_t - \phi)\Big) + \frac{Y_t}{r}\Big(V_a(1) - V_a(\pi_t)\Big) + \frac{1}{r}\frac{dV_a(\pi_t)}{dt}. \tag{3.3}$$

The law of motion of π_t implies that when $Y_t \neq 0$,

$$\frac{dV_a(\pi_t)}{dt} = -\pi_t(1 - \pi_t)Y_t V_a'(\pi_t).$$

Next, suppose the agent deviates and chooses a_t different from the market's believed effort $\hat{a}_t$, his continuation payoff at π_t is:[19]

$$\Big(\pi_t \hat{a}_t - c(a_t - \phi)\Big) + \frac{Y_t}{r}\Big(V_a(1) - V_a(\pi_t)\Big) - \frac{\hat{Y}_t}{r}\pi_t(1 - \pi_t)V_a'(\pi_t), \tag{3.4}$$

with $\hat{Y}_t \equiv \mu\chi_t \hat{a}_t$. This is because the agent's effort is unobservable, deviations only affect his cost of effort as well as the arrival rate of breakthroughs. But conditional on no disclosure, π_t is still updated according to the '*believed effort*'. Choosing a_t to maximize (3.4) obtains the agent's best response correspondence:

$$a_t \begin{cases} = \phi & \text{when } \frac{\mu\chi_t}{r}(V_a(1) - V_a(\pi_t)) < c \\ \in [\phi, 1] & \text{when } \frac{\mu\chi_t}{r}(V_a(1) - V_a(\pi_t)) = c \\ = 1 & \text{when } \frac{\mu\chi_t}{r}(V_a(1) - V_a(\pi_t)) > c \end{cases}. \tag{3.5}$$

The term $\frac{\mu\chi_t}{r}(V_a(1) - V_a(\pi_t))$ measures his '*reputational incentive*', which is increasing in the publicity of his performance ($\mu\chi_t$); decreasing in the discount rate (r) and increasing in the difference between his

[18] The following ODE uses the fact that in equilibrium, the market's belief is always correct, i.e.$\hat{a}_t = a_t$.

[19] Expression (3.4) deals with the case in which $\hat{Y}_t \neq 0$. When $\hat{Y}_t = 0$, the agent's continuation payoff following a deviation is $(\pi_t \hat{a}_t - c(a_t - \phi)) + \frac{Y_t}{r}(V_a(1) - V_a(\pi_t))$.

continuation value when $\pi_t = 1$ and his current continuation value. After a breakthrough is disclosed, $\pi_t = 1$ and the agent's reputational incentive disappears. So $a_t = w_t = V_a(1) = \phi$.[20] Since χ_t is bounded below 1, he has an incentive to work only if his continuation value is low enough, i.e.

$$V_a(\pi_t) \leq \overline{V}_a \equiv \phi - \frac{cr}{\mu}. \tag{3.6}$$

The Intermediary's Incentives: Let $V_m(\pi_t)$ be the intermediary's continuation value conditional on $\theta = 1$. Similar to $V_a(\pi_t)$, $V_m(\pi_t)$ satisfies the following ODE:

$$V_m(\pi_t) = a_t(1 - \pi_t) + \frac{Y_t}{r}\Big(b - V_m(\pi_t)\Big) + \frac{1}{r}\frac{dV_m(\pi_t)}{dt}. \tag{3.7}$$

Since her deviations cannot be observed by the market either, her best response correspondence is:

$$\chi_t \begin{cases} = 1 & \text{when } V_m(\pi_t) < b \\ \in [0,1] & \text{when } V_m(\pi_t) = b \\ = 0 & \text{when } V_m(\pi_t) > b \end{cases}. \tag{3.8}$$

From (3.8), the intermediary is more reluctant to disclose information when $V_m(\pi_t)$ is high and vice versa. Intuitively, due to the lack of commitment, her incentive to disclose information only depends on the revenue she can milk from the agent by continuing the relationship after she knows that $\theta = 1$.

Patience Level Conditions: As in other dynamic game models, the discount rate r matters for the equilibrium outcomes. Therefore, I introduce the following patience level condition:

Condition 1. *Players' patience level is high if $r < \frac{\mu\phi}{c}(1 - c)$ and is low otherwise.*

Patience level is high is equivalent to $\overline{V}_a > c\phi$. This condition is less demanding when the talent premium ϕ is higher, the arrival rate μ is higher or the cost of effort c is lower.

Indifference & Value Invariance Curves: I define two curves, which are of critical importance in analyzing the long run players' dynamic incentives. First, for a given time interval (t_0, t_1), suppose the intermediary is indifferent between disclosing and withholding information for all $t \in (t_0, t_1)$, then

[20] By replacing $V_a(1)$ with any other constant V^*, it is straightforward to extend my result to the case in which the agent receives continuation value V^* after a breakthrough is disclosed, which is assumed in Bonatti and Hörner (2015).

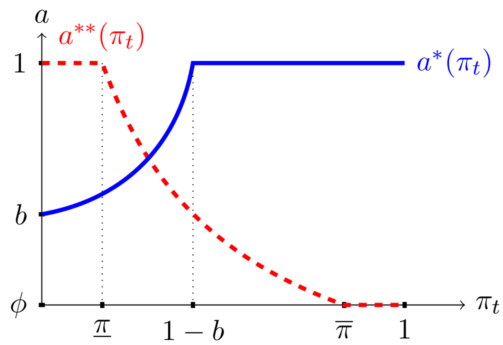

Figure 1: Indifference Curve (solid) and Value Invariance Curve (dashed)

$V_m(\pi_t) = b$ and $\frac{dV_m(\pi_t)}{dt} = 0$. Equation (3.7) implies that $a_t(1-\pi_t) = b$. Let

$$a^*(\pi_t) \equiv \min\left\{1, \frac{b}{1-\pi_t}\right\} \tag{3.9}$$

be the intermediary's '*indifference curve*', which is increasing in π_t and becomes flat once π_t exceeds $1-b$.

Second, when patience level is high i.e. $\overline{V}_a > c\phi$, let $a^{**}(\pi_t)$ be defined as:

$$a^{**}(\pi_t) \equiv \begin{cases} 1 & \text{when } \pi_t \leq \underline{\pi} \\ \frac{\phi(1-c)-\frac{rc}{\mu}}{\pi_t} & \text{when } \pi_t \in (\underline{\pi}, \overline{\pi}] \\ \phi & \text{when } \pi_t > \overline{\pi}, \end{cases} \tag{3.10}$$

in which $\underline{\pi} \equiv \phi(1-c) - \frac{cr}{\mu}$ and $\overline{\pi} \equiv (1-c) - \frac{cr}{\mu\phi}$. The high patience level condition ensures that $0 < \underline{\pi} < \overline{\pi} < 1$. I call $a^{**}(\pi_t)$ the agent's '*value invariance curve*', which is flat at both ends and is strictly decreasing when $\pi_t \in [\underline{\pi}, \overline{\pi}]$. Intuitively, if $\pi_t \in [\underline{\pi}, \overline{\pi}]$, $V_a(\pi_t) = \overline{V}_a$ and $\chi(\pi_t) = 1$, then the agent's continuation value remains unchanged (i.e. $\frac{dV_a(\pi_t)}{dt} = 0$) if $a(\pi_t) = a^{**}(\pi_t)$. I depict the two curves together in Figure 1.

3.2 Exogenous Information Benchmark

In this subsection, I consider the benchmark scenario in which $\chi_t = 1$ for all t, i.e. the market directly observes every breakthrough the agent has achieved.[21] I call this the '*exogenous information*' benchmark, which will later be compared with the strategic disclosure case. Proposition 1 characterizes the unique

[21]Several alternative interpretations of this benchmark scenario include: when the agent can directly communicate with the market, or when the intermediary is required to disclose all information.

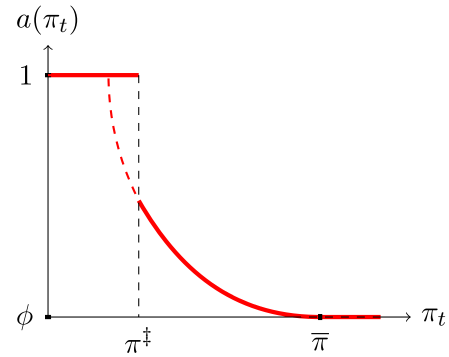

Figure 2: Effort Path under Exogenous Information (High Patience)

MPE in this benchmark:

Proposition 1 (Exogenous Information)**.** *There exists a unique MPE. When patience level is low, $a(\pi_t) = \phi$ for all $\pi_t \in (0, 1]$. When patience level is high, there exists $\pi^\ddagger \in (\underline{\pi}, 1)$ such that:*

$$a(\pi_t) = \begin{cases} 1 & \text{when } \pi_t \leq \pi^\ddagger \\ a^{**}(\pi_t) & \text{when } \pi_t > \pi^\ddagger. \end{cases}$$

The agent's effort path when patience level is high is shown in Figure 2. Since this benchmark scenario fits into the definition of '*Poisson good news model*' in Faingold and Sannikov (2011), the unique MPE is also the unique Nash Equilibrium. According to Proposition 1, the agent's effort is increasing over time when r is low enough. As in other Poisson good news models, for example, Board and Meyer-ter-Vehn (2013), the agent works too little when π_t is high (relative to social first best), leading to '*procrastination inefficiencies*'. This is because under exogenous information, the agent's marginal benefit from exerting effort is $\frac{\mu}{r}(\phi - V_a(\pi_t))$ and $V_a(\pi_t)$ is non-decreasing in π_t. The cutoff belief, $\pi^\ddagger$, is pinned down by the '*promise keeping condition*', i.e. his continuation value at $\pi^\ddagger$ is $\overline{V}_a$. Effort is interior when $\pi_t \in (\pi^\ddagger, \overline{\pi})$ since the agent's incentive to exert effort is decreasing with the market's '*believed effort*'. This uniquely pins down the equilibrium effort level.

3.3 Equilibrium Characterization

I characterize the unique MPE when information is strategically disclosed by an intermediary. The main result is stated as Proposition 2.

Proposition 2. *This game admits a unique MPE.*

- *When patience level is low, $\chi(\pi_t) = 1$ and $a(\pi_t) = \phi$ for all $\pi_t \in (0, 1]$.*

- *When patience level is high, there exists a cutoff belief $\pi^\dagger \in (\underline{\pi}, 1)$, such that:*

$$a(\pi_t) \equiv \begin{cases} a^{**}(\pi_t) & \text{when } \pi_t > \pi^\dagger \\ a^{*}(\pi_t) & \text{when } \pi_t \leq \pi^\dagger \end{cases}, \quad \chi(\pi_t) \equiv \begin{cases} 1 & \text{when } \pi_t > \min\{1-b, \pi^\dagger\} \\ \frac{cr}{\mu(\phi - V_a(\pi_t))} & \text{when } \pi_t \leq \min\{1-b, \pi^\dagger\} \end{cases}.$$

Moreover, $a^*(\pi^\dagger) > a^{**}(\pi^\dagger)$ and $\chi(\pi_t)$ is strictly increasing in π_t when $\pi_t \leq \min\{1-b, \pi^\dagger\}$.

From now on, I will focus on the high patience case, due to the presence of non-trivial reputation building behaviour, which is economically interesting. Effort and disclosure rate when patience level is high are depicted in Figures 3 and 4 as functions of π_t, depending on whether $\pi^\dagger$ is below (interior case) or above (corner case) $1-b$. Common in both figures, effort is inverse U-shaped while the rate of disclosure decreases over time. Moreover, the equilibrium effort path coincides with the intermediary's indifference curve when the market is pessimistic and coincides with the agent's value invariance curve otherwise.

For some intuition, since the intermediary benefits from extracting the agent's effort, she has more incentive to disclose information when future effort is lower and vice versa. When patience level is low, the agent has no incentive to exert effort and the intermediary cannot milk much revenue from the agent. Because of this, she strictly prefers to disclose information. When patience level is high, the agent has more incentive to work when π_t is lower, as in the exogenous information benchmark. Anticipating this, the intermediary will have an incentive to withhold information when π_t is low.

This gives rise to the following countervailing incentive structure: the agent has more incentive to exert effort when his performance is more visible while the intermediary has more incentive to disclose information when the agent shirks. When the intermediary is required to stochastically withhold information, effort and disclosure rate must be chosen to make the other player indifferent, which results in the described dynamic pattern. The cutoff belief $\pi^\dagger$ is pinned down by the '*promise keeping condition*': the agent's continuation value at $\pi^\dagger$ must be $\overline{V}_a$, the highest continuation value under which he can be motivated to exert effort.

Mapping this back into the application to professional service industries, my result speaks for the anecdotal evidence that firms tend to establish their star employees early on in their careers. But for those who fail to succeed in the beginning, their latter successes will be discounted more and more heavily and are less visible to the outside market. This distinction is driven by the employer's strategic motives: the worker's eagerness to build up his reputation enables his employer to milk more revenue from his hard work, which explains why he is not being promoted for a long time. This lack of transparency in performance will frustrate the worker in the long run and his effort eventually decreases.

The predictions of my model in terms of wages (w_t) and promotion rates ($a_t\chi_t$) match the empirical

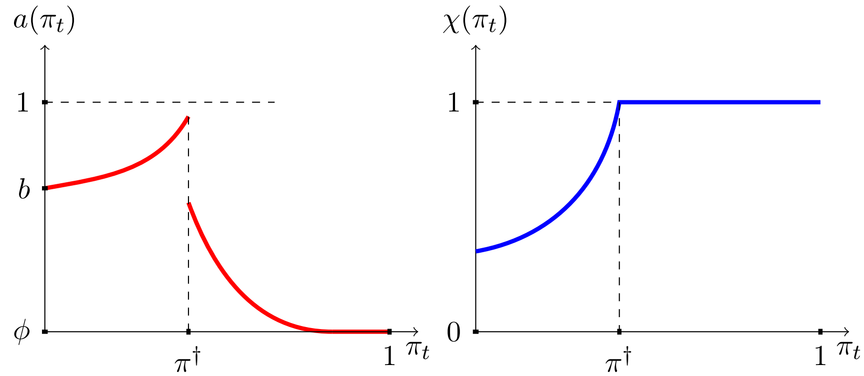

Figure 3: Effort and Disclosure Dynamics in the unique MPE: High Patience & Interior Case ($\pi^\dagger \leq 1-b$)

findings in Baker, Gibbs and Holmström (1994a,b), who study the wage and promotion dynamics empirically using 20 years of personnel data from a large US firm.[22] First, the real wage of a worker conditional on remaining at the entry level job (level 2 in their paper) is first increasing and then decreasing as a function of time (Figure IV on page 951). Second, the promotion rate from level 2 to level 3 is also an inverse U-shaped function of tenure (the number of years at level 2), which is documented in their Table IV (page 902). I provide a novel explanation for these facts based on the dynamic interactions between the agent's incentive to build up his reputation and the intermediary's incentive to release information to the labour market.[23]

Why MPE: The unique MPE is robust to private monitoring, i.e. when the agent cannot perfectly observe the arrival of breakthroughs. Private monitoring is prevalent in applications where performance evaluations are subjective (Fuchs [2007]). In my model, MPE is attractive since players' strategies only depend on the public belief and their incentives do not rely on the fine details of their private histories. As a result, their incentive constraints are still satisfied even when breakthroughs are based on the intermediary's subjective assessment.[24]

Furthermore, this unique MPE remains to be an equilibrium when the intermediary can disclose past

[22]The predictions of my model, that effort is inverse U-shaped and the disclosure rate is decreasing are robust when the market can also learn from infrequently arrived public signals, when the intermediary's benefit from disclosure is low, i.e. $b \in (0, \phi]$, when the intermediary has different time preferences (for example, having a different discount factor or is finitely lived), when the agent does not know θ, etc. These will be discussed in section 6. They are also robust when we consider other PBEs that are non-Markov (section 4).

[23]Alternative explanations in the existing literature include, for example, the inverse U-shaped promotion rate can also be explained by a combination of market learning and human capital accumulation.

[24]The robustness of equilibria under private monitoring has been discussed extensively in the repeated games literature, for example, Mailath and Morris (2002), etc. The gist of this literature is: coordinating current play on past histories becomes more challenging when monitoring is private and histories are not commonly known among players.

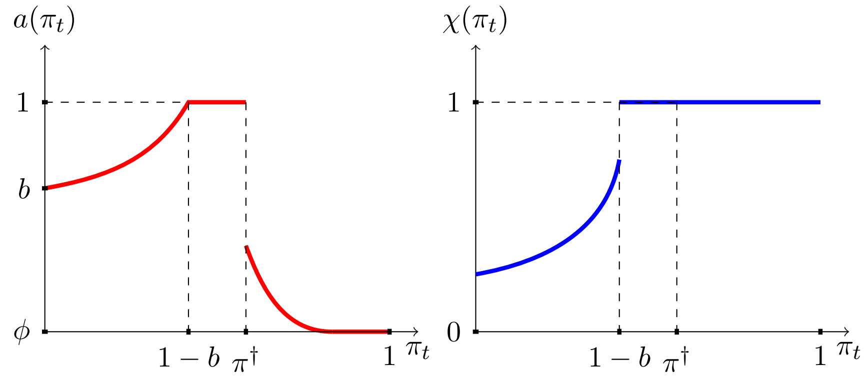

Figure 4: Effort and Disclosure Dynamics in the unique MPE: High Patience & Corner Case ($\pi^\dagger > 1-b$)

breakthroughs. This is because her incentive to withhold information is weakly increasing over time, i.e. given she has an incentive to withhold information today, she also has an incentive to withhold it in the future. Nonetheless, the possibility of disclosing past breakthroughs introduces additional equilibria, which will be discussed in section 6.

Remark: My result that effort is inverse U-shaped is different from, albeit complementary to, both career concern models with Gaussian learning and reputation building models with Poisson good news.

In Holmström (1999), the agent's effort is decreasing over time, which is driven by the Gaussian information structure, i.e. the market's signal is highly sensitive to effort. As time elapses, the precision of market belief increases and the impact of current effort decreases.[25] Contrary to Holmström (1999), I examine cases in which the market receives information *infrequently*. This is relevant to reputation building of junior people, as the market rarely knows about their performance before they become famous.

In Poisson good news models, for example, Board and Meyer-ter-Vehn (2013), the agent's effort is increasing over time as the market's belief deteriorates since $\phi - V_a(\pi_t)$ is decreasing in π_t. In my model, the agent's effort is inverse U-shaped.[26] This is driven by the intermediary's incentive to release information: as the market becomes more pessimistic, the intermediary can capture a larger share of the agent's surplus once she knows that $\theta = 1$. Effort is decreasing over time along the intermediary's

[25]In a recent paper, Hörner and Lambert (2015) examine a variant of the Holmström model in which the agent's talent is changing over time, following a mean reverting process. In these '*changing type*' models, the agent's effort level can be constant over time, which is the case in their stationary equilibria.

[26]Inverse U-shaped effort is also reported in Bonatti and Hörner (2015), although the driving forces are very different. In their model, the agent does not know his type, as in the strategic experimentation literature. His effort decreases as his confidence in his own ability falls. In my model, declining effort is driven by the intermediary's incentives to suppress information, which comes from the strategic interaction between the intermediary and the agent, not through the agent's private learning.

indifference curve in order to motivate her to stochastically disclose information.

3.4 Proof of Proposition 2

I prove the result in seven steps, with proofs of the lemmas relegated to Appendix A (high patience case) and Online Appendix B (low patience case). To approach the problem, I use the observation that changes in w_t vanishes when π_t goes to 0, that is, the agent's problem becomes '*approximately stationary*' at the limit. I pin down his limiting continuation value when $\pi_t \to 0$ and characterize players' limiting equilibrium behaviours. The agent's continuation value at low enough π_t is characterized by a limiting value problem, which admits a unique solution. This solution is then used to compute $\pi^\dagger$.

Step 1: Market Learning First, I show that the market eventually learns about the agent's type. This result validates my approach of analyzing players' values and behaviours at the limiting belief.

Lemma 3.1. *In every MPE, $\pi_t(1-\pi_t)$ converges to 0 in probability.*[27]

An implication is that $\chi_t \neq 0$ for all t and hence, $Y_t > 0$. Therefore, we can re-write (3.3) as:

$$V_a(\pi_t) = \Big(\pi_t a_t - c(a_t - \phi)\Big) + \frac{Y_t}{r}\Big(V_a(1) - V_a(\pi_t) - \pi_t(1-\pi_t)V_a'(\pi_t)\Big).$$

Step 2: Continuation Value & Behaviour at the Limiting Belief Let

$$V_a(0) \equiv \begin{cases} c\phi & \text{when } r < \frac{\mu\phi(1-c)}{c} \\ \frac{\mu\phi^2}{r+\mu\phi} & \text{when } r \geq \frac{\mu\phi(1-c)}{c} \end{cases}, \quad V_m(0) \equiv \begin{cases} b & \text{when } r < \frac{\mu\phi(1-c)}{c} \\ \frac{r\phi+\mu b}{r+\mu} & \text{when } r \geq \frac{\mu\phi(1-c)}{c} \end{cases}.$$

Lemma 3.2 shows that the limits of players' continuation values, $\lim_{\pi_t\to 0} V_a(\pi_t)$ and $\lim_{\pi_t\to 0} V_m(\pi_t)$, exist. It also explicitly computes these values as well as players' equilibrium behaviours at the limit:

Lemma 3.2. *Players' limiting continuation values exist. Furthermore, in every MPE,*

$$\lim_{\pi_t\to 0} V_a(\pi_t) = V_a(0) \text{ and } \lim_{\pi_t\to 0} V_m(\pi_t) = V_m(0),$$

and there exists $\varepsilon > 0$ such that for all $\pi_t \in (0,\varepsilon)$,

- *$\chi(\pi_t) = 1$ and $a(\pi_t) = \phi$ when patience level is low.*

[27] In my model, the unique MPE also coincides with the unique public equilibrium, i.e. players' strategies depend only on the public history. Aside from π_t, the market also observes calendar time. In the Online Appendix, I show $\chi(\pi_t) > 0$ under a weaker assumption that players' strategies are public, which implies the existence of a 1-to-1 mapping between t and π_t. As a result, t does not contain any extra information in addition to π_t.

- *$a(\pi_t) = a^*(\pi_t)$ and $\chi(\pi_t) = \frac{cr}{\mu(\phi - V_a(\pi_t))}$ when patience level is high.*

In what follows, I will focus on the high patience case. Lemma 3.2 implies that when patience level is high and π_t is close to 0, the intermediary must be indifferent between disclosing and withholding information and the agent's effort must be $a^*(\pi_t)$. Hence, the agent's continuation value when π_t is low enough is characterized by a solution to the following limiting value problem:

$$V_a(\pi_t) = \Big(\pi_t a^*(\pi_t) - c(a^*(\pi_t) - \phi)\Big) + \frac{\mu\chi(\pi_t)a^*(\pi_t)}{r}\Big(\phi - V_a(\pi_t) - \pi_t(1-\pi_t)V_a'(\pi_t)\Big) \tag{3.11}$$

with $\lim_{\pi_t \to 0} V_a(\pi_t) = c\phi$. Accordingly, $\chi(\pi_t)$ must be chosen to provide him an incentive to choose an interior effort level, i.e.

$$\frac{\mu\chi(\pi_t)}{r}\Big(\phi - V_a(\pi_t)\Big) = c. \tag{3.12}$$

Plugging (3.12) into (3.11) results in the following ODE:

$$V_a(\pi_t) = \Big(\pi_t a^*(\pi_t) + c\phi\Big) - \frac{ca^*(\pi_t)\pi_t(1-\pi_t)V_a'(\pi_t)}{\phi - V_a(\pi_t)}, \tag{3.13}$$

with $\lim_{\pi_t \to 0} V_a(\pi_t) = c\phi$. Lemma 3.3 shows that this problem admits a unique solution:

Lemma 3.3. *The above limiting value problem admits a unique solution $V_a^*(\pi_t)$. When patience level is high, there exists $\varepsilon > 0$ such that $V_a(\pi_t) = V_a^*(\pi_t)$ for all $\pi_t \in (0, \varepsilon)$.*[28]

The proof uses the observation that different values of $V_a(\cdot)$ at any interior belief will lead to diverging values when $\pi_t \to 0$. As a result, the limiting value uniquely pins down the value at every interior belief.

Step 3: Value Invariance Curve The next Lemma exploits the implications of the agent's value invariance curve.

Lemma 3.4. *Suppose $\chi_t = 1$ and $V_a(\pi_t) = \overline{V}_a$,*

- *If $\pi_t < \underline{\pi}$, then $V_a'(\pi_t) < 0$ for all $a_t \in [\phi, 1]$. If $\pi_t > \overline{\pi}$, then $V_a'(\pi_t) > 0$ for all $a_t \in [\phi, 1]$.*
- *If $\pi_t \in [\underline{\pi}, \overline{\pi}]$, then*

$$V_a'(\pi_t) \begin{cases} > 0 & \text{when } a_t > a^{**}(\pi_t) \\ = 0 & \text{when } a_t = a^{**}(\pi_t) \\ < 0 & \text{when } a_t < a^{**}(\pi_t). \end{cases}$$

[28] ODE (3.13) can be transformed into a Bernoulli Equation, by letting $Z(\pi_t) \equiv \phi - V_a(\pi_t)$. This admits a closed form solution in limiting form. However, this formula is inconvenient both for characterizing the equilibrium as well as for doing comparative statics. Hence, throughout my analysis, I use an indirect approach to establish the uniqueness as well as other properties of MPE.

The proof is straightforward from (3.3), which is omitted. Intuitively, fixing $V_a(\pi_t)$, higher believed effort in $(t-dt, t]$ leads to higher wages in $(t-dt, t]$ and therefore, higher $V_a(\pi_{t-dt}) - V_a(\pi_t)$. Since π_t is decreasing over time, this implies that $V_a'(\pi_t)$ is larger.

Step 4: Constructing an MPE I construct an MPE in the high patience case and later show its uniqueness. Let

$$\pi^\dagger \equiv \sup\left\{\pi_t \middle| V_a^*(\pi) < \overline{V}_a \text{ for all } \pi \in (0, \pi_t)\right\}. \tag{3.14}$$

By definition, if $\pi^\dagger < 1$,[29] then $V_a(\pi^\dagger) = \overline{V}_a$ and $\lim_{\pi_t \uparrow \pi^\dagger} V_a'(\pi_t) > 0$. I claim that the strategy profile displayed in Proposition 2 and its induced conditional belief constitute an MPE. This also implies that the agent's continuation value when $\pi_t < \pi^\dagger$ is $V_a^*(\pi_t)$, which is the unique solution to the limiting value problem in Lemma 3.3.

I check players' incentive constraints. The agent's incentive constraints are satisfied when $\pi_t \leq \min\{1-b, \pi^\dagger\}$ since the choice of $\chi(\pi_t)$ makes him indifferent between working and shirking. When $\pi_t \in (\min\{1-b, \pi^\dagger\}, \pi^\dagger]$, the definition of $\pi^\dagger$ implies that $V_a(\pi_t) \leq \overline{V}_a$ and the intermediary fully discloses information, implying that he has a strict incentive to exert effort. When $\pi_t > \pi^\dagger$,

- If $a^{**}(\pi_t) > \phi$, then his continuation value remains $\overline{V}_a$ and he is indifferent.
- If $a^{**}(\pi_t) = \phi$, then his continuation value is weakly above $\overline{V}_a$ and he weakly prefers to shirk.

The intermediary's incentive constraint when $\pi_t \leq \min\{1-b, \pi^\dagger\}$ is satisfied since her flow payoff is constantly b thereafter. To verify her incentive constraints when $\pi_t > \min\{1-b, \pi^\dagger\}$, I only need to show:

$$a^*(\pi^\dagger) > a^{**}(\pi^\dagger). \tag{3.15}$$

Since $\lim_{\pi_t \uparrow \pi^\dagger} V_a'(\pi_t) > 0$, (3.15) is then implied by Lemma 3.4.

Step 5: Monotonicity of Disclosure Rates I show that $\chi(\pi_t)$ is strictly increasing in π_t for all $\pi_t \leq \min\{1-b, \pi^\dagger\}$. It is equivalent to show that $V_a^*(\pi_t)$ is strictly increasing in π_t, or equivalently, the agent's equilibrium continuation value is decreasing over time, conditional on no disclosure.

Lemma 3.5. *When $\pi_t \leq \pi^\dagger$, $V_a^*(\pi_t)$ is strictly increasing in π_t.*

[29] I will defer the proof of $\pi^\dagger < 1$ to Lemma 3.8, the proof of which does not rely on any of the previous Lemma.

Step 6: Uniqueness of MPE First, I show that if $\pi_t \le \min\{1-b, \pi^\dagger\}$, then $a(\pi_t) = a^*(\pi_t)$ and $\chi(\pi_t)$ must be chosen to make the agent indifferent. Let

$$\pi^1 \equiv \sup\left\{\pi_t \middle| \pi_t \in (0, \min\{1-b, \pi^\dagger\}],\ a(\pi) = a^*(\pi) \text{ for all } \pi < \pi_t\right\}.$$

I show the following Lemma:

Lemma 3.6. *In every MPE, $\pi^1 = \min\{1-b, \pi^\dagger\}$.*

Next, I show that the intermediary always fully discloses information when $\pi_t > \min\{1-b, \pi^\dagger\}$. If $\pi^\dagger > 1-b$ and $\pi_t \in (1-b, \pi^\dagger)$, then $V_a^*(\pi_t) \le \overline{V}_a$ and the intermediary always strictly prefers to disclose, so $a(\pi_t) = \chi(\pi_t) = 1$. When $\pi_t > \pi^\dagger$, Lemma 3.7 shows that effort must be $a^{**}(\pi_t)$ in every MPE. Let

$$\pi^2 \equiv \sup\left\{\pi_t \middle| \pi_t \in [\pi^\dagger, 1) \text{ and } a(\pi) = a^{**}(\pi) \text{ for all } \pi \in [\pi^\dagger, \pi_t)\right\}.$$

Lemma 3.7. *In every MPE, $\pi^2 = 1$.*

Step 7: Range of $\pi^\dagger$ It is already known from Step 4 that $a^{**}(\pi^\dagger) < a^*(\pi^\dagger)$, which gives a strictly positive lower bound on $\pi^\dagger$, i.e. the lowest belief at which $a^*(\cdot)$ and $a^{**}(\cdot)$ intersect. In this step, I establish an upper bound:

$$\pi^\dagger \le 1 - \frac{cr}{\mu\phi}. \tag{3.16}$$

I show the following claim which implies (3.16): when the agent's reputation is π_t, given that the market's belief is correct, his continuation value is at least $\pi_t\phi$ for any Markov disclosure plan. This minimum is achieved when $\chi(\pi_t) = 0$ and $a(\pi_t) = \phi$, i.e. no information is disclosed and market's belief is π_t forever. This result is also interesting by itself since it characterizes the harshest possible punishment to the agent when the intermediary can commit to disclosure plans.

Formally, for all $\mathbf{a}, \hat{\mathbf{a}} \in \mathbf{A}$, $\boldsymbol{\chi} \in \mathbf{X}$ and $\pi_t \in (0,1)$, let $\Pi^{\hat{\mathbf{a}},\boldsymbol{\chi}}(\pi_t, \mathbf{a})$ be the agent's continuation payoff when he adopts effort plan $\mathbf{a}$ under disclosure rule $\boldsymbol{\chi}$, believed effort process $\hat{\mathbf{a}}$ and initial reputation π_t. Let

$$V^{\hat{\mathbf{a}},\boldsymbol{\chi}}(\pi_t) \equiv \max_{\mathbf{a}\in\mathbf{A}} \Pi^{\hat{\mathbf{a}},\boldsymbol{\chi}}(\pi_t, \mathbf{a})$$

be his continuation value. Consider the following program, which characterizes his lowest continuation value under the restriction that $\hat{\mathbf{a}}$ is optimal for the agent under $\hat{\mathbf{a}}$ and $\boldsymbol{\chi}$:

$$\mathcal{V}(\pi_t) \equiv \min_{(\hat{\mathbf{a}},\boldsymbol{\chi})\in \mathbf{A}\times\mathbf{X}} V^{\hat{\mathbf{a}},\boldsymbol{\chi}}(\pi_t), \tag{3.17}$$

subject to:

$$V^{\hat{\mathbf{a}},\boldsymbol{\chi}}(\pi_t) = \Pi^{\hat{\mathbf{a}},\boldsymbol{\chi}}(\pi_t, \hat{\mathbf{a}}).$$

Lemma 3.8. *When patience level is high, $\mathcal{V}(\pi_t) = \pi_t\phi$ for all $\pi_t \in (0,1)$,*

As a direct implication, $V_a(\pi_t) > \overline{V}_a$ when $\pi_t > 1 - \frac{cr}{\mu\phi}$, which implies that $\pi^\dagger \leq 1 - \frac{cr}{\mu\phi}$. This result is not trivial since when $\chi(\pi_t) > 0$, higher effort level becomes incentive compatible. When the market anticipates this, π_t also declines faster conditional on no disclosure.[30] As a result, playing the best response to $\chi(\pi_t) = 0$ and $\hat{a}(\pi_t) = \phi$, which is $a(\pi_t) = \phi$, cannot guarantee him payoff $\pi_t\phi$.

3.5 Comparisons Between Exogenous and Endogenous Information

I compare the effort path in the unique MPE with the equilibrium effort path under the exogenous information benchmark to assess the impact of strategic information disclosure. First, effort is inverse U-shaped under endogenous information as opposed to monotone increasing under exogenous information. This is because effort needs to be low enough in order to provide the intermediary an incentive to disclose information. Second, the cutoffs at which effort jumps up, $\pi^\dagger$ and $\pi^\ddagger$, are also different.

The next result compares $\pi^\dagger$ with $\pi^\ddagger$. To see why this is interesting, notice that the two effort paths (as well as the paths of disclosure rates) coincide when $\pi_t > \max\{\pi^\dagger, \pi^\ddagger\}$. Moreover, effort is strictly lower under endogenous information when π_t is small enough. The remaining question is: whether strategic disclosure can lead to higher effort at some intermediate beliefs. This is only the case when $\pi^\dagger > \pi^\ddagger$. If so, the presence of a strategic intermediary motivates the agent to front-load effort.

Intuitively, one would expect that front-loading will always happen since χ_t is decreasing over time, that is, the intermediary withholds information at a higher rate at more pessimistic beliefs. This reduces the agent's continuation value at optimistic beliefs. Anticipating this, he will work harder early on since if he shirks today, his performance will be less visible tomorrow.

Unfortunately, the above logic is flawed, since it ignores the impact of equilibrium disclosure rate on the speed of market learning, which affects the dynamics of wages. To see this, when c is very close

[30] This effect is also discussed in Cisternas (2015) using a career concern model, in which the agent has more incentive to exert higher effort when the market anticipates higher effort, albeit his equilibrium payoff can be lower.

to 0, χ_t is also very close to 0 even for fairly high π_t, so is $\mu\chi_t a_t$, the rate of market learning. If this is the case, the market will attribute the absence of news to the intermediary's low disclosure rate, instead of the agent's incompetence. As a result, the agent's flow payoff, $a_t\pi_t - c(a_t - \phi)$, will remain high for a long time. In contrast, π_t deteriorates much faster when information is exogenous. So the agent's short run payoff is higher under endogenous information albeit his long run payoff is lower. Since $r > 0$, for c small enough, there exists an interior belief such that the agent's continuation value is higher under endogenous information when π_t exceeds this belief and vice versa. In this case, having a strategic intermediary censoring information can exacerbate the procrastination problem. This effect is more pronounced when r and μ are high, c and ϕ are low. Proposition 3 provides a sufficient condition:

Proposition 3. *For every r and ϕ, there exist $\overline{c} > 0$ and $\underline{\mu} > 0$ satisfying*

$$\phi - \overline{c}\phi - \frac{r\overline{c}}{\underline{\mu}} > 0$$

such that for every $c < \overline{c}$ and $\mu > \underline{\mu}$, there exists an open subset $B \subset (\phi, 1)$, such that $\pi^{\ddagger} > \pi^{\dagger}$ when $b \in B$.

To summarize, the unique MPE exhibits two sources of inefficiencies. First, the agent's effort needs to be low enough to encourage the intermediary to disclose information, and the latter is necessary to sustain reputation building incentives. Second, withholding information when $\pi_t < \pi^{\dagger}$ does not necessarily lower the agent's continuation value: when c is low and withholding information happens too early, it is actually encouraging the agent to procrastinate more.

Remark: It is worth clarifying that Proposition 3 is about the incentives to front-load effort instead of welfare. Welfare comparisons are hard to obtain due to the inconvenience of expressing the agent's equilibrium continuation value and the intermediary's equilibrium disclosure rate explicitly, which makes computing social surplus not tractable. Even when $\pi^{\dagger} > \pi^{\ddagger}$, despite effort is front-loaded, there is no guarantee that endogenous disclosure can improve social welfare. This is because first, welfare depends on the prior belief π_0 at which to evaluate payoffs; and second, welfare depends not only on the timing of effort, but also on its entire dynamics.

4 Semi-Markov Equilibria

In this section, I show that allowing players' strategies to condition on additional payoff irrelevant state variables can mitigate the inefficiencies of the unique MPE. In particular, I focus on a solution concept that minimally departs from MPE, Semi-Markov Equilibrium (SME, see Definition 3), which allows the intermediary and the agent to coordinate on payoff irrelevant variables off the equilibrium path.

I characterize the set of SME outcomes, which turns out to be tractable and contains both the unique MPE and the exogenous information equilibrium as special cases. To characterize this set, I construct a subclass of SMEs that covers the entire set of SME outcomes.

Proposition 4. *For every $\pi^{\S} \in [0, \min\{1-b, \pi^{\dagger}\}]$, there exist π^* and $\pi^{**} \in [\pi^{\S}, 1)$, such that the following 'three phase strategy profile' forms an SME.*

- ***Phase I:*** *If $\pi_t \leq \pi^{\S}$, then $a(\pi_t) = a^*(\pi_t)$ and $\chi(\pi_t)$ is chosen to make the agent indifferent.*
- ***Phase II:*** *If $\pi_t > \pi^{\S}$ and the intermediary has never concealed a breakthrough in the past, then $\chi(\pi_t) = 1$ and*
$$a(\pi_t) \equiv \begin{cases} 1 & \text{when } \pi_t \in (\pi^{\S}, \pi^*] \\ a^{**}(\pi_t) & \text{when } \pi_t > \pi^*. \end{cases}$$
- ***Phase III:*** *If $\pi_t > \pi^{\S}$ and the intermediary has concealed a breakthrough in the past, then*
$$a(\pi_t) \equiv \begin{cases} a^*(\pi_t) & \text{when } \pi_t \in (\pi^{\S}, \pi^{**}] \\ \phi & \text{when } \pi_t > \pi^{**}. \end{cases}$$
$\chi_t = 1$ when $\pi_t > \min\{1-b, \pi^{**}\}$ and χ_t is chosen to make the agent indifferent when $\pi_t \leq \min\{1-b, \pi^{**}\}$.

Moreover, every SME is outcome equivalent to an SME described above.[31]

Figure 5 depicts the effort and disclosure dynamics in a generic three phase strategy profile. For some intuition, SME allows players' strategies to condition on another state variable, that is, whether the intermediary has concealed breakthroughs in the past or not. Every SME is characterized by a cutoff belief, $\pi^{\S}$, below which withholding information happens on the equilibrium path. When $\pi_t \leq \pi^{\S}$ (Phase I), the equilibrium play resembles that in the unique MPE in which effort and disclosure rate are chosen to make both players indifferent. When $\pi_t > \pi^{\S}$ and the intermediary has never deviated before (Phase II), the agent exerts high effort and the intermediary fully discloses information. Suppose the intermediary has deviated before (Phase III), players coordinate on the low effort low disclosure rate equilibrium. π^* and π^{**} are chosen such that the agent's on-path continuation value at π^* and off-path continuation value at π^{**} are both $\overline{V}_a$. A formal characterization of the pair (π^*, π^{**}) is presented in Appendix C, along with the proof of Proposition 4.

[31] The set of three phase strategy profiles characterize all possible on-path behaviours in SMEs. However, there exist other SMEs which differ in terms of their off-path play.

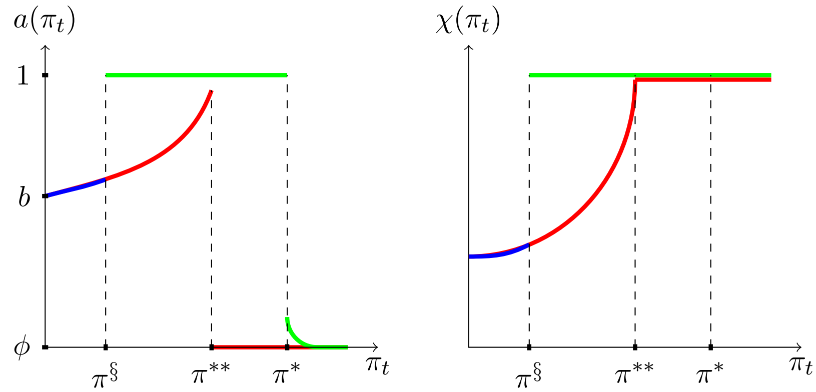

Figure 5: Effort and disclosure rate in an SME (in three phase strategy profile) with cutoff $\pi^{\S}$, with Phase I in blue, Phase II in green and Phase III in red.

Mapping back into the applications to law and consulting industries, SMEs capture the following effect which is absent in the unique MPE: the agent becomes frustrated if his past successes have been ignored and in response to that, his future effort will decrease. Suppressing information today leads to lower effort in the future. This is reflected in the equilibrium behaviour that effort jumps downward once a breakthrough is concealed.[32]

In what follows, I discuss the properties of SMEs. First, notice that the unique MPE is the SME with the highest $\pi^{\S}$, i.e. $\pi^{\S} = \min\{1-b, \pi^{\dagger}\}$, while the unique equilibrium in the exogenous information benchmark is outcome equivalent to the SME with $\pi^{\S} = 0$. SME allows $\pi^{\S}$ to take any value between 0 and $\min\{1-b, \pi^{\dagger}\}$, which enables players to control the timing at which withholding information starts and implies an efficiency gain when the optimal cutoff does not coincide with the unique MPE cutoff. In particular, when c is low, players can coordinate on a low enough $\pi^{\S}$ so that the market learns rapidly when π_t is high. This circumvents the problem identified in Proposition 3 since the agent cannot sustain high flow payoff at any interior belief. Moreover, SME boosts effort to 1 when $\pi_t \in (\pi^{\S}, \min\{1-b, \pi^{\dagger}\}]$. This arrangement is incentive compatible since the intermediary has no decision to make before the first breakthrough arrives, so increasing the agent's effort before that does not upset her incentive constraints. Some other properties of SME are summarized below.

Front-loading Effort: Since the unique MPE and the equilibrium under exogenous information are both special cases of SME, the optimal SME must (weakly) outperform both.

[32] To comment more on the jump, effort jumps from 1 to $a^*(\pi_t)$ when $\pi_t \in (\pi^{\S}, \pi^{**}]$; from 1 to ϕ when $\pi_t \in (\pi^{**}, \pi^*]$; from $a^{**}(\pi_t)$ to ϕ when $\pi_t \in (\pi^*, \overline{\pi}]$.

Corollary 4.1. *There exists $\varepsilon > 0$ such that for all $\pi^{\S} \in (0, \varepsilon)$, $\pi^*(\pi^{\S}) > \pi^{\ddagger}$.*

Corollary 4.1 says that SME with $\pi^{\S}$ small enough strictly outperforms the unique equilibrium in the exogenous information benchmark in terms of front-loading effort, i.e. the agent starts working harder at a higher belief.

Simple Belief Updating Rule: SMEs have the attractive property that the market does not need to compute the *probability that a breakthrough has been withheld in the past* in order to formulate its posterior belief about θ. This is because in every SME, the market always correctly anticipates the agent's on-path effort and the intermediary's on-path disclosure rate.[33]

Robustness: Every SME induced by a three phase strategy profile remains robust when the intermediary can disclose past breakthroughs. This is because her incentive to withhold information is weakly increasing over time, both on and off the equilibrium path. However, SMEs (aside from the unique MPE) are not robust to private monitoring, since they rely on the arrival times of breakthroughs being common knowledge between the intermediary and the agent.

5 Optimal Markov Policy with Commitment

In this section, I examine the optimal Markov policy when the intermediary can commit to dynamic disclosure plans: the intermediary designs a Markov policy, $(\mathbf{a}, \boldsymbol{\chi}) \in \mathbf{A} \times \mathbf{X}$, to maximize her expected payoff evaluated at belief π_0, subject to the market's and the agent's incentive constraints.[34]

Recall the definitions of $V^{\mathbf{a},\boldsymbol{\chi}}(\pi)$ and $\Pi^{\mathbf{a},\boldsymbol{\chi}}(\pi, \mathbf{a})$ in subsection 3.4. The intermediary's maximization problem is:

$$\sup_{(\mathbf{a},\boldsymbol{\chi})\in\mathbf{A}\times\mathbf{X}} \Big\{\pi_0 \underbrace{\Big(r\int_0^\infty e^{-rt-(y_t-y_0)}(a_t - w_t - b)dt + b\Big)}_{\text{gain from high type}} - (1-\pi_0)\underbrace{r\int_0^\infty e^{-rt} w_t dt}_{\text{loss from low type}}\Big\}, \tag{5.1}$$

[33]Nevertheless, the set of PBEs in which players' strategies can condition on π_t as well as whether breakthrough has been withheld in the past or not is much larger than the set of SMEs, since in general, future effort and disclosure rate can differ whenever the intermediary has withheld a breakthrough in the past, regardless of whether play has gone off-path or not. However, complicated market belief updating and wage formulas will occur, which are not tractable to analyze.

[34]It is worth emphasizing that there is a loss of generality when focusing on Markov Policies. This is because potentially, effort and disclosure rate can also depend on players' private histories as well as calendar time, on top of the public belief.

subject to:

$$\underbrace{V^{\mathbf{a},\chi}(\pi_0) = \Pi^{\mathbf{a},\chi}(\pi_0, \mathbf{a})}_{\text{Effort plan } \mathbf{a} \text{ is optimal for the agent}}, \; \dot{\pi}_t = -\pi_t(1-\pi_t)\mu a(\pi_t)\chi(\pi_t) \text{ and } w(\pi_t) = \pi_t a(\pi_t),$$

with $y_t \equiv \ln \frac{1-\pi_t}{\pi_t}$. Re-write the intermediary's objective function as:

$$r(1-\pi_0)\int_0^\infty e^{-rt-y_t}(a_t - w_t - b)dt - r(1-\pi_0)\int_0^\infty e^{-rt}w_t dt + b\pi_0. \tag{5.2}$$

Plugging in the expression for w_t and ignoring positive affine transformations, the intermediary's problem is to minimize:

$$\int_0^\infty e^{-rt}\frac{\pi_t}{1-\pi_t}dt. \tag{5.3}$$

subject to the agent's and the market's incentive constraints. The optimal policy is characterized below:

Proposition 5. *An optimal Markov policy exists. For every $\pi_0 \in (0,1)$, there exists $0 \leq \pi' < \pi'' \leq \pi_0$ such that the optimal Markov policy with commitment has three phases:*

- ***Shirking Phase:*** *If $\pi_t > \pi''$, then $\chi(\pi_t) = 1$ and $a(\pi_t) = \phi$.*
- ***Working Phase:*** *If $\pi_t \in (\pi', \pi'']$, then $a(\pi_t) = \chi(\pi_t) = 1$.*
- ***Deadline:*** *If $\pi_t \leq \pi'$, then $\chi(\pi_t) = 0$ and $a(\pi_t) = \phi$.*

Moreover, $\pi' > 0$ if and only if $\pi_0 > \pi^\ddagger$.

The proof uses standard optimal control techniques, which is relegated to Online Appendix C. As long as π_0 is sufficiently large, the shirking phase is not degenerate ($\pi'' < \pi_0$) and the deadline is not trivial ($\pi' > 0$). The agent's effort is still inverse U-shaped, with no effort when the market's belief is extremely high or low and maximal effort when the market's belief is intermediate.

Comparing with the unique MPE or the SMEs, the optimal commitment solution has three interesting features. First, the intermediary commits to fully disclose information when π_t is high. Second, the agent's effort is always bang-bang. Third, when the prior π_0 is high enough, she commits to a deadline. This reduces the agent's continuation value and makes $a_t = 1$ incentive compatible at more optimistic beliefs.[35]

These features contrast sharply with the equilibria without commitment. Aside from the exogenous information equilibrium ($\pi^\S = 0$), information disclosure rate and effort must both be interior when π_t is low enough in every other SME. Moreover, the third feature (deadline) is not replicable in any *PBE* in absence of commitment, as formally stated in the following Corollary:

[35] As shown in Lemma 3.8, deadline is the harshest possible punishment.

Corollary 5.1. *In every PBE, there exists no* $h_m^t \in H_m$*, such that* $\chi(h_m^{t'}) = 0$ *for all* $h_m^{t'} \succeq h_m^t$*.*

This is because the agent knows the intermediary's private history, and he will shirk forever after reaching h_m^t. Anticipating this, the intermediary is tempted to disclose information at h_m^t. Lack of commitment unravels the deadline arrangement, making it harder to motivate the agent when π_t is high.

6 Discussions & Extensions

In this section, I enrich and extend the baseline model in several directions and examine the robustness of my results. In subsection 6.1, I discuss the properties of Markov Equilibria when the intermediary can disclose past breakthroughs. In subsection 6.2, I allow the market to learn from a public signal in addition to the intermediary's private signal. In subsection 6.3, I discuss several variants of the model, by investigating different preferences of the intermediary and the agent, alternative informational assumptions, etc.

6.1 Disclosing Past Breakthroughs

In this subsection, I allow the intermediary to disclose past breakthroughs. Despite every SME (including the unique MPE) remains robust to disclosing past breakthroughs, enriching the intermediary's strategy space opens up new equilibrium possibilities. This is true even for Markov solution concepts, since the possibility of disclosing past breakthroughs changes the set of payoff relevant state variables.

Formally, when disclosing past breakthroughs is allowed, I adopt the '*multi-stage game*' formulation in Murto and Välimäki (2013). In what follows, I will proceed at an intuitive level and will formally define the game in Online Appendix D.2. Let $x_t \in \{0, 1\}$ be defined as:

$$x_t \equiv \begin{cases} 1 & \text{when } h_m^t \neq \{\varnothing\} \\ 0 & \text{when } h_m^t = \{\varnothing\}, \end{cases}$$

which determines the intermediary's capability of disclosing information. Once $x_t = 1$, whether the breakthrough arrives at t or before is payoff irrelevant. The market's belief about θ as well as x_t are both deterministic functions of t. Hence, t and x_t are the only payoff relevant state variables.

The solution concept is weak Markov Perfect Equilibrium (or wMPE), in which at time t, the agent's effort depends only on t and x_{t-}, but the intermediary's disclosure rate depends not only on t and x_{t-}, but also on $\xi_t \equiv \mathbf{1}\{t \in h_m^t\}$, i.e. whether a signal has arrived at t or not.[36] By definition, every SME constructed in section 4 (therefore, also the unique MPE) is a wMPE. Despite the set of wMPEs is much larger than the set of SMEs and is not tractable to analyze, the next Lemma identifies some common properties of wMPEs:

Lemma 6.1. *In every wMPE, there exists $\pi^{\P} \in [0, 1)$, such that:*

- *The intermediary always fully discloses the first breakthrough when $\pi_t > \pi^{\P}$.*
- *The agent's effort is $a^*(\pi_t)$ after the first breakthrough is concealed when $\pi_t \leq \pi^{\P}$.*

There are multiple equilibria even under a given $\pi^{\P}$ since effort and disclosure rate are not uniquely pinned down when $\pi_t \leq \pi^{\P}$ but before a breakthrough has been concealed.

Although the possibility of disclosing past breakthroughs generates additional equilibria, in which the rate of learning, the cutoff beliefs and the effort paths are *quantitatively* different from the SMEs in the baseline model, two *qualitative* features remain robust. First, the intermediary only withholds information when π_t is low. Second, if π_t falls below the cutoff and a breakthrough has been concealed, the agent's effort is decreasing over time and coincides with $a^*(\pi_t)$ in *every* wMPE.

6.2 Learning from Public Signals

In this subsection, the market can also learn from public signals, arrive according to Poisson rate $\mu_0 \theta a_t$, in addition to the private signals disclosed by the intermediary, which arrive at Poisson rate $\mu_1 \theta a_t$, with parameters $\mu_0 \geq 0, \mu_1 > 0$. The market automatically observes the public signal, but can only observe the private signal after the intermediary discloses it. Let $\mu \equiv \mu_0 + \mu_1$ be the '*net arrival rate*', I introduce the following condition, which measures the public signal arrival rate.

Condition 2. *Public signal arrival rate is low if $\mu_0 \leq \frac{cr}{\phi(1-c)}$.*

Otherwise, we say that public signal arrival rate is high. Intuitively, the intermediary has less control over the market's information if public signal arrival rate is high and vice versa. The patience level condition is re-defined as follows:

Condition 3. *Patience level is high if $r < \frac{\mu\phi}{c}(1 - c)$. Patience level is low if $r \geq \frac{\mu\phi}{c}(1 - c)$.*

[36] If we insist on more restrictive solution concepts, for example, by requiring that the agent's effort to depend only on t and x_{t-} and the intermediary's disclosure rate to depend only on t and x_t, only a trivial equilibrium exists, in which the intermediary always discloses information. I will state and show this result in Online Appendix D.2.2.

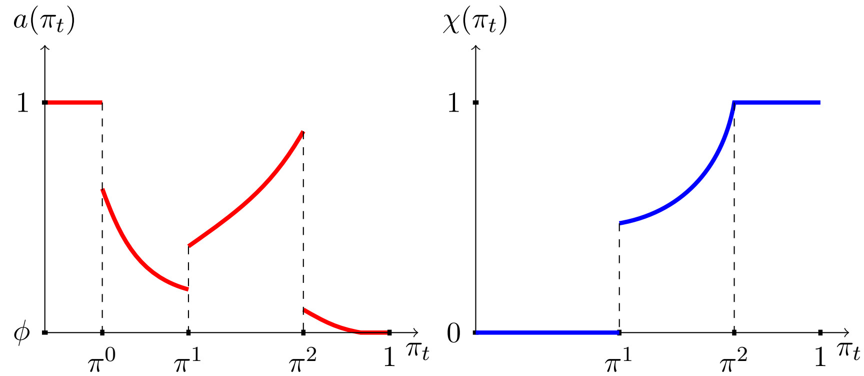

Figure 6: Effort and Disclosure Rate when Public Signals Arrive Frequently

Since $\mu > \mu_0$, patience level must be high when public signal arrival rate is high. I start with the low arrival rate case, in which the unique MPE and the set of SME outcomes coincide with those in the baseline model.

Lemma 6.2. *When* $\mu_0 \leq \frac{cr}{\phi(1-c)}$,

- *The unique MPE with public signals is outcome equivalent to the unique MPE in the baseline model.*
- *Every SME with public signals is outcome equivalent to an SME in the baseline model.*

Intuitively, fixing the net arrival rate μ and given that μ_0 is small enough, the exact decomposition between public and private signals has no impact on the equilibrium outcome since the net disclosure rate $\mu_0 + \mu_1\chi_t$ is unchanged: we can always find $\chi_t \in [0, 1]$ to match the equilibrium net disclosure rate in the baseline model ($\mu_0 = 0$). In this case, public and private signals are perfect substitutes.

Once μ_0 exceeds $\frac{cr}{\phi(1-c)}$, there exists no χ_t that can match the net disclosure rate when π_t is small. Due to frequent arrival of public signals, the agent has a strict incentive to exert effort when π_t is low even if the intermediary ceases to disclose private signals. As a result, $\chi_t \to 0$ and $a_t \to 1$ when π_t is close to 0. Moreover, public and private signals are no longer perfect substitutes: an increase in μ_0 increases the agent's continuation value at some beliefs, which leads to an increase in the net disclosure rate required to motivate the agent, so the disclosure rate of private signal can increase with μ_0 through this indirect intertemporal effect. The equilibrium effort path and disclosure rate are shown in Figure 6.[37]

6.3 Other Extensions & Discussions

In this subsection, I discuss several alternative specifications of the intermediary's and the agent's payoffs, the robustness of my result to alternative informational assumptions as well as how to enrich the baseline

[37]The proofs of which are available upon request.

model to account for more realistic features.

The Intermediary's Time Preference: The effort and disclosure dynamics in every SME has nothing to do with the intermediary's discount factor. As a result, it remains robust when the intermediary faces a different discount rate. The unique MPE is robust even when the intermediary is finitely lived. This is because whenever she is supposed to partially disclose information, her flow payoff is b from then on; whenever she is supposed to fully disclose information, her flow payoff is weakly below b at every instant.[38]

The Intermediary's Disclosure Benefit is Low: When $b \in (0, \phi]$,[39] withholding information is the intermediary's dominant strategy when π_t is sufficiently low. The agent ceases to exert effort and market learning stops when π_t falls below the following cutoff: $\pi^* \equiv 1 - \frac{b}{\phi}$, at which the intermediary's continuation payoff is b and the agent's continuation value is $\pi_t \phi$. However, in the unique MPE, effort is still inverse U-shaped and the disclosure rate is still decreasing over time (both reaching 0 when $\pi_t \leq \pi^*$), which are qualitatively similar to the baseline model.

The Intermediary as a Supervisor: In academia and professional sports, the intermediary who has private information about the junior agent's performance is usually his direct supervisor, instead of his current employer. Different from employers, although supervisors gain private benefits from extracting her agent's effort and can obtain a network benefit from establishing her agent in front of the public, they do not pay their agent's wages. As a result, their flow payoff when $\pi_t < 1$ is θa_t, instead of $\theta a_t - w_t$.

When patience level is high, the game still admits a unique MPE, in which the effort and disclosure rate dynamics are shown in Figure 7, with the red dashed line being the agent's value invariance curve (same as the baseline model) and the blue dashed line being the intermediary's indifference curve. Similar to the baseline model, the intermediary's disclosure rate is decreasing over time and the agent's effort coincides with the value invariance curve when $\pi_t > \pi^\dagger$ and coincides with the intermediary's indifference curve when $\pi_t \leq \pi^\dagger$. However, since the intermediary only cares about the agent's effort, her indifference curve is flat. As a result, the agent's effort remains unchanged when belief is low.

Convex Effort Cost: When the agent's effort cost is convex instead of linear, the baseline model still admits a unique MPE. Several features of the unique MPE under linear cost remain robust, including

[38] All SMEs in three phase strategy profiles constructed in section 4 remain robust to finitely lived intermediaries if the next intermediary inherits all her predecessors' information.

[39] Cases in which $b \geq 1$ and $b < 0$ are trivial, since the intermediary always has a strict incentive to disclose or to withhold information

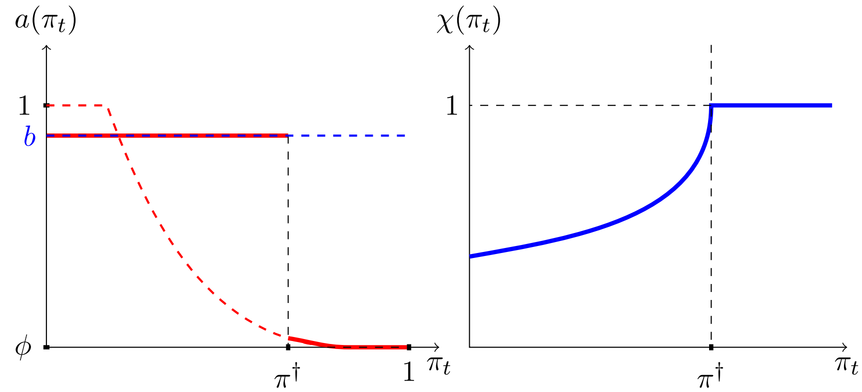

Figure 7: Effort and Disclosure Rate when the Intermediary is the Agent's Supervisor

inverse U-shaped effort and decreasing disclosure rate. However, there will be no '*jump*' in effort under convex cost, which is the main difference from the linear cost model.[40]

Comments on Informational Assumptions: The qualitative features of the effort and disclosure dynamics identified in my baseline model are robust to other variants of informational assumptions. For example, even if the intermediary can observe effort or the agent cannot observe the breakthrough, the unique MPE remains robust. When the agent does not know his type, as in career concern models (for example, Holmström [1999]), the agent's equilibrium effort can condition not only on the market's belief, but also on his private belief as well as the market's belief about his private belief, all of these become payoff relevant.[41] Despite characterizing the set of MPEs is a formidable task, but nonetheless, the agent's effort will still be inverse U-shaped. To see this, if the agent knows his type, his effort path will be the same as in the '*reputation concern case*'. If he does not know his type, his effort is also inverse U-shaped, as shown in Bonatti and Hörner (2015). This is because the agent will become more pessimistic about his ability as no breakthrough has arrived over time. Due to the complementarity between effort and ability, he has less incentive to exert effort. To summarize, we would still anticipate inverse U-shaped effort and decreasing disclosure rate.

Multiple Job Levels: My model can be enriched to account for the fact that lots of '*promising future stars*' (especially in academia and professional sports) work hard despite having favourable public beliefs. To see this, suppose there are three job levels: $1, 2$ and 3, and three types: high, medium and low.

[40] However, there will be kinks in the equilibrium effort path.

[41] The intermediary's private belief is trivial, since she can only disclose information after her private belief is 1.

The high type and the medium type can produce breakthroughs in level 1 while only the high type can produce breakthroughs in level 2. After a breakthrough in level k is disclosed, the intermediary in level k receives a lump sum payoff and the agent is promoted to level $k+1$, working for a new intermediary.

In this variation of my model, high type agents have stronger incentives to work hard in level 1 for two reasons.[42] First, he can be further promoted after reaching level 2, thus his continuation value of entering level 2 is higher than the medium type. Second, being promoted to level 2 at an earlier date distinguishes himself from the medium type, leading to a higher market belief and increases the intermediary's information disclosure rate at level 2. As a result, high type agents work harder in level 1. In terms of predictions, those who have succeeded earlier in level 1 are more likely to be promoted sooner in level 2, which is the well-known '*fast track*' phenomena documented in the personnel economics literature. Moreover, '*fast tracks*' are more salient when information is disclosed by a strategic intermediary, due to the decreasing disclosure rate over time.

7 Conclusion

This paper studies the impact of endogenous information disclosure by a strategic intermediary on an agent's incentive to build up his reputation. In the unique MPE, the agent's effort is inverse U-shaped and the information disclosure rate is decreasing over time. Surprisingly, the agent's continuation value can be higher and can have more incentives to procrastinate in the unique MPE comparing with the exogenous information benchmark. This is because withholding information also slows down the rate of market learning, which allows the agent to enjoy high flow payoff for a long time, even without producing any breakthroughs. Relaxing the Markov restriction and allowing players to coordinate on payoff irrelevant events such as whether breakthroughs have been withheld in the past or not, can mitigate this inefficiency. This is because it enables players to flexibly choose the cutoff belief, below which the intermediary starts to withhold information.

[42] As in the baseline model, the agent will eventually stop working hard once reaching the final level, i.e. level 3.

References

[1] Baker, George, Michael Gibbs and Bengt Holmström (1994a) "The Internal Economics of the Firm: Evidence from Personnel Data," *Quarterly Journal of Economics*, 109(4), 881-919.

[2] Baker, George, Michael Gibbs and Bengt Holmström (1994b) "The Wage Policy of a Firm," *Quarterly Journal of Economics*, 109(4), 920-955.

[3] Bergemann, Dirk and Ulrich Hege (1998) "Venture Capital Financing, Moral Hazard, and Learning," *Journal of Banking & Finance*, 22, 703-735.

[4] Bergemann, Dirk and Ulrich Hege (2005) "The Financing of Innovation: Learning and Stopping." *RAND Journal of Economics*, 36(4), 719-752.

[5] Board, Simon and Moritz Meyer-ter-Vehn (2013) "Reputation for Quality," *Econometrica*, 81(6), 2381-2462.

[6] Bonatti, Alessandro and Johannes Hörner (2015) "Career Concerns with Exponential Learning," *Theoretical Economics*, forthcoming.

[7] Campbell, Arthur, Florian Ederer and Johannes Spinnewijn (2014) "Delay and Deadlines: Freeriding and Information Revelation in Partnerships," *American Economic Journal: Microeconomics*, 6(2), 163-204.

[8] Che, Yeon-Koo and Johannes Hörner (2015) "Optimal Design for Social Learning," Working Paper, Colimbia University and Yale University.

[9] Cisternas, Gonzalo (2015) "Two-Sided Learning and Moral Hazard," Working Paper, MIT Sloan.

[10] Dewatripont, Mathias, Ian Jewitt and Jean Tirole (1999) "The Economics of Career Concerns, Part II: Application to Missions and Accountability of Government Agencies," *Review of Economic Studies*, 66(1), 199-217.

[11] Ekmekci, Mehmet (2011) "Sustainable Reputations with Rating Systems," *Journal of Economic Theory*, 146(2), 479-503.

[12] Ely, Jeffrey (2015) "Beeps," Working Paper, Northwestern University.

[13] Faingold, Eduardo and Yuliy Sannikov (2011) "Reputation in Continuous-Time Games," *Econometrica*, 79(3), 773-876.

[14] Fuchs, William (2007) "Contracting with Repeated Moral Hazard and Private Evaluations," *American Economic Review*, 97(4), 1432-1448.

[15] Halac, Marina, Navin Kartik and Qingmin Liu (2015) "Contests for Experimentation," Working Paper, Columbia University.

[16] Halac, Marina and Andrea Prat (2015) "Managerial Attention and Worker Engagement," Working Paper, Columbia University.

[17] Holmström, Bengt (1999) "Managerial Incentive Problems: A Dynamic Perspective," *Review of Economic Studies*, 66(1), 169-182.

[18] Hörner, Johannes and Nicholas Lambert (2015) "Motivational Ratings," Working Paper, Yale University and Stanford University.

[19] Hörner, Johannes and Larry Samuelson (2013) "Incentives for Experimenting Agents," *RAND Journal of Economics*, 44(4), 632-663.

[20] Kahn, Lisa, B. (2013) "Asymmetric Information between Employers," *American Economic Journal: Applied Economics*, 5(4), 165-205.

[21] Keller, Godfrey, Sven Rady and Martin Cripps (2005) "Strategic Experimentation with Exponential Bandits," *Econometrica*, 73(1), 39-68.

[22] Klein, Nicolas and Sven Rady (2011) "Negatively Correlated Bandits," *Review of Economic Studies*, 78(2), 693-732.

[23] Kremer, Ilan, Yishay Mansour and Motty Perry (2015) "Implementing the *Wisdom of Crowd*", *Journal of Political Economy*, 122(5), 998-1012.

[24] Mailath, George and Stephen Morris (2002) "Repeated Games with Almost-Public Monitoring," *Journal of Economic Theory*, 102, 189-228.

[25] Mailath, George and Larry Samuelson (2001) "Who Wants a Good Reputation?" *Review of Economic Studies*, 68(2), 415-441.

[26] Maskin, Eric and Jean Tirole (2001) "Markov Perfect Equilibrium I: Observable Actions," *Journal of Economic Theory*, 100, 191-219.

[27] Milgrom, Paul and Sharon Oster (1987) "Job Discrimination, Market Forces and the Invisibility Hypothesis," *Quarterly Journal of Economics*, 102(3), 453-476.

[28] Milgrom, Paul (2008) "What the Seller Won't Tell You: Persuasion and Disclosure in Markets," *Journal of Economic Perspectives*, 22(2), 115-131.

[29] Murto, Pauli and Juuso Välimäki (2013) "Delay and Information Aggregation in Stopping Games with Private Information," *Journal of Economic Theory*, 148, 2404-2435.

[30] Waldman, Michael (1984) "Job Assignments, Signalling, and Efficiency," *RAND Journal of Economics*, 15(2), 255-267.

A Proof of Proposition 2

In this Appendix, I show the Lemmas that lead to Proposition 2.

A.1 Proof of Lemma 3.1

Suppose towards a contradiction, that there exists $\pi^* \in (0,1)$ such that $\lim_{t\to\infty} \pi_t = \pi^*$, where π_t is the market's belief conditional on $h^t = \{\varnothing\}$. Rewrite (3.1) as:

$$\frac{d\pi_t}{\pi_t(1-\pi_t)} = -Y_t dt.$$

Integrate both sides from 0 to ∞, we have:

$$\int_0^\infty Y_t dt = \ln \frac{1-\pi^*}{\pi^*} - \ln \frac{1-\pi_0}{\pi_0}, \tag{A.1}$$

Since $\pi^*, \pi_0 \in (0,1)$, so $\int_0^\infty Y_t dt$ is finite. This implies that $\lim_{t\to\infty} Y_t = 0$. Since $Y_t = \mu\chi_t a_t$, $\mu > 0$ and $a_t \geq \phi$, we have $\chi_t \to 0$. Hence, for every $\varepsilon > 0$, there exists $T \in \mathbb{R}_+$ such that $\chi_t < \varepsilon$ for all $t > T$. Pick ε such that:

$$\frac{\mu\varepsilon}{r}\phi < c. \tag{A.2}$$

Since $V_a(\pi_t) \geq 0$, so (3.5) implies that $a_t = \phi$ for all $t > T$, implying that $V_m(\pi_t) < b$. This suggests that $\chi_t = 1$ for all $t > T$, which leads to a contradiction.

A.2 Proof of Lemma 3.2

Since $\chi(\pi_t) > 0$ for all t, let $\hat{V}_a(t) \equiv V_a(\pi_t)$. Replace $V_a(\pi_t)$ and $V_a'(\pi_t)$ with $\hat{V}_a(t)$ and $\hat{V}_a'(t)$ in (3.3), we have:

$$\hat{V}_a(t) = \pi_t a_t + c\phi + a_t\Big(\frac{\mu\chi_t}{r}(\phi - \hat{V}_a(t)) - c\Big) + \frac{1}{r}\hat{V}_a'(t).$$

The unique bounded solution is:

$$\hat{V}_a(t) = r\int_t^\infty e^{-r(s-t)}\Big[\pi_s a_s + c\phi + a_s\Big(\frac{\mu\chi_s}{r}(\phi - \hat{V}_a(s)) - c\Big)\Big]ds = c\phi + X_t + W_t, \tag{A.3}$$

where

$$X_t \equiv r\int_t^\infty e^{-r(s-t)} Z_s ds, \quad Z_t \equiv a_t\Big(\frac{\mu\chi_t}{r}(\phi - \hat{V}_a(t)) - c\Big),$$

and

$$W_t \equiv r\int_t^\infty e^{-r(s-t)}\pi_s a_s ds.$$

This implies that for every $\varepsilon > 0$, there exists $T \in \mathbb{R}^+$ such that for every $t > T$:

$$\pi_t < \varepsilon, \; W_t < \varepsilon,$$

The expression for $\hat{V}_a(t)$ suggests that:

$$X_t < \hat{V}_a(t) - c\phi < X_t + \varepsilon. \tag{A.4}$$

Moreover, according to the agent's incentive constraint, Z_t has the following property:

$$Z_t \begin{cases} < 0 & \text{when } \frac{\mu\chi_t}{r}(\phi - V_a(\pi_t)) < c \\ = 0 & \text{when } \frac{\mu\chi_t}{r}(\phi - V_a(\pi_t)) = c \\ > 0 & \text{when } \frac{\mu\chi_t}{r}(\phi - V_a(\pi_t)) > c \end{cases}$$

The following proof focuses on $t \in \mathbb{R}^+$ large enough such that $\pi_t < 1 - b$, i.e. $a^*(\pi_t) \in (\phi, 1)$. Recall the patience level is high if and only if:

$$r < \frac{\mu\phi(1-c)}{c} \text{ or } \overline{V}_a > c\phi.$$

The following Lemma claims that the agent's continuation value converges in the limit.

Lemma A.1. *There exists $V_a(0) \in \mathbb{R}_+$ such that:*[43]

$$\lim_{t\to\infty} \hat{V}_a(t) = \lim_{\pi_t \to 0} V_a(\pi_t) = \lim_{\pi_t \to 0} V_a(\pi_t).$$

The proof of this Lemma is in Section B of the Online Appendix. Let $V_a(0) \equiv \lim_{t\to\infty} \hat{V}_a(t)$. In what follows, I will show that $V_a(0) = c\phi$ and then characterize players' behaviours in the limit.

A.2.1 Limiting Continuation Value in High Patience Case

I show that $V_a(0) = c\phi$ by ruling out all other possibilities.

Part I: Suppose $V_a(0) > \overline{V}_a$. Since $V_a(0) > \overline{V}_a > c\phi$, there exists $T \in \mathbb{R}^+$ such that $\hat{V}_a(t) > \overline{V}_a$ and $X_t > 0$ for all $t > T$. But

$$\frac{\mu\chi_t}{r}(\phi - V_a(\pi_t)) < c$$

for all $\chi_t \in [0,1]$ when $V_a(\pi_t) > \overline{V}_a$, and hence, $Z_t < 0$ for all $t > T$. This leads to a contradiction.

Part II: Suppose $V_a(0) < c\phi$. Then there exists $T \in \mathbb{R}^+$ such that $\hat{V}_a(t) < \overline{V}_a$ and $X_t < 0$ for all $t > T$. Hence, there exists $t > T$ such that $Z_t < 0$, i.e.

$$a(\pi_t) = \phi, \quad \frac{\mu\chi_t}{r}\left(\phi - \hat{V}_a(t)\right) < c.$$

Since $\hat{V}_a(t) < \overline{V}_a$, so $\chi_t \in (0,1)$, implying that $V_m(\pi_t) \geq b$. Admissibility requires the existence of $\varepsilon_0 > 0$ such that

$$a(\pi_{t+\varepsilon_1}) \in \left[\phi, a^*(\pi_{t+\varepsilon_1})\right) \text{ for all } \varepsilon_1 \in (0, \varepsilon_0).$$

Equation (3.7) implies that $V_m(\pi_{t+\varepsilon_1}) > b$, thus $\chi_{t+\varepsilon_1} = 0$, contradicting the conclusion of Lemma 3.1.

[43]The fact that $\lim_{t\to\infty} \hat{V}_a(t) = \lim_{\pi_t\to 0} V_a(\pi_t)$ is a direct implication of Lemma 3.1.

Part III: Suppose $V_a(0) \in (c\phi, \overline{V}_a)$. Then there exists $T \in \mathbb{R}^+$ such that $\hat{V}_a(t) < \overline{V}_a$ and $X_t > 0$ for all $t > T$. This implies the existence of $t > T$ such that:

$$a(\pi_t) = 1, \quad \frac{\mu \chi_t}{r}\Big(\phi - \hat{V}_a(t)\Big) > c.$$

The intermediary's incentive constraint requires that $\chi(\pi_t) > 0$ and $V_m(\pi_t) \leq b$. Admissibility (left-continuity with respect to belief or equivalently, right-continuity with respect to time) requires the existence of $\varepsilon_0 > 0$ such that

$$a(\pi_{t+\varepsilon_1}) \in \big(a^*(\pi_{t+\varepsilon_1}), 1\big) \text{ for all } \varepsilon_1 \in (0, \varepsilon_0).$$

Equation (3.7) implies that $V_m(\pi_{t+\varepsilon_1}) < b$ for all $\varepsilon_1 \in (0, \varepsilon_0)$, which further implies that $\chi_{t+\varepsilon_1} = 1$ for all $\varepsilon_1 \in (0, \varepsilon_0)$. But because $\hat{V}_a(t) < \overline{V}_a$ for all $t > T$, so $a_t = 1$ when $\chi_t = 1$. But then since $a^*(\pi_t) < 1$ for all $t > T$, we have $V_m(\pi_t) > b$, contradicting the previous conclusion that $V_m(\pi_t) \leq b$.

Part IV: Suppose $V_a(0) = \overline{V}_a$. Then there exists $T \in \mathbb{R}^+$ such that $X_t > 0$ for all $t > T$. If $V_a(\pi_t) < \overline{V}_a$ for all $\pi_t \in (0, \pi_T)$, we can obtain the same contradiction as in Part III.

Suppose there exists $t > T$ such that $\hat{V}_a(t) \geq \overline{V}_a$. Since $X_t > 0$ and $Z_t \leq 0$ for all t such that $\hat{V}_a(t) \geq \overline{V}_a$, so there exists $t' > t$ with $V_a(\pi_{t'}) < \overline{V}_a$. Since $V_a(\cdot)$ is a continuous function, there exists $t'' \in [t, t')$ such that

$$\hat{V}_a(t'') = \overline{V}_a, \quad \hat{V}_a'(t'') < 0.$$

Since $\hat{V}_a(t'') = \overline{V}_a$ suggests that $Z_{t''} \leq 0$, we have:

$$\overline{V}_a = c\phi + \underbrace{Z_{t''} + \frac{1}{r}\hat{V}_a'(t'')}_{<0} < c\phi,$$

which is a contradiction.

A.2.2 Behaviour at Limiting Belief

Next, I characterize players' equilibrium behaviours when π_t is close to 0. Since $V_a(0) = c\phi < \overline{V}_a$, there exists $\nu > 0$ such that $V_a(\pi_t) < \overline{V}_a$ for all $\pi_t < \nu$. The entire discussion will be focusing on $\pi_t < \nu$.

I start with showing that $a(\pi_t) < a^*(\pi_t)$ cannot happen when π_t is small enough. Suppose towards a contradiction, that there exists $\pi_t < \nu$ such that $a(\pi_t) < a^*(\pi_t)$, then since $a(\cdot)$ is left continuous, there exists $\varepsilon > 0$ such that $a(\pi') < a^*(\pi')$ for all $\pi' \in (\pi_t - \varepsilon, \pi_t]$. Since $\chi \neq 0$ for all t, and $V_a(\pi') < \overline{V}_a$, it has to be the case that $\chi(\pi') \in (0, 1)$, implying that $V_m(\pi') = b$ for all $\pi' \in (\pi_t - \varepsilon, \pi_t]$, and hence $V_m'(\pi') = 0$ for all $\pi' \in (\pi_t - \varepsilon, \pi_t)$. From (3.7), we have $a(\pi') = a^*(\pi')$ for all $\pi' \in (\pi_t - \varepsilon, \pi_t)$. But this implies that $a(\pi_t) = a^*(\pi_t)$, which is a contradiction.

Then I show that $a(\pi_t) > a^*(\pi_t)$ also cannot happen when $\pi_t \to 0$. The previous step implies the existence of $\nu > 0$ such that $a(\pi_t) \geq a^*(\pi_t)$ for all $\pi_t < \nu$, implying that $V_m(\pi_t) \geq b$. Since $\chi(\pi_t) > 0$, so $V_m(\pi_t) = b$, which is achieved only when $a(\pi_t) = a^*(\pi_t)$ for all $\pi_t \leq \nu$. This completes the proof.

A.3 Proof of Lemma 3.3

The existence and uniqueness of solution to the initial value problem is established in Lemma A.1 and the existence of solution to the limiting value problem is established in Lemma A.2 in the Online Appendix. To show uniqueness in the limiting value problem, let $Z(\pi_t) \equiv \phi - V_a(\pi_t)$, re-write (3.13) as

$$-Z(\pi_t)^2 + Z(\pi_t)\Big(\phi - c\phi - \pi_t a^*(\pi_t)\Big) = ca^*(\pi_t)\pi_t(1 - \pi_t)Z'(\pi_t). \tag{A.5}$$

Suppose towards a contradiction that there exists two solutions, $Z_1(\pi_t) \neq Z_2(\pi_t)$. Since $\lim_{\pi_t \to 0} Z_1(\pi_t) = \lim_{\pi_t \to 0} Z_2(\pi_t) = \phi - c\phi$, there exists π^* small enough such that:

$$Z_1(\pi), Z_2(\pi) \in [\phi - c\phi - \varepsilon, \phi - c\phi + \varepsilon], \text{ for all } \pi < \pi^*$$

and

$$Z_1(\pi^*) = Z_1 \neq Z_2 = Z_2(\pi^*).$$

where $\varepsilon \in (0, \frac{\phi - c\phi}{2})$. Without loss of generality, let $Z_1 > Z_2$. I show that $Z_1(\pi_t) - Z_2(\pi_t) \geq Z_1 - Z_2 > 0$ for all $\pi_t \in (0, \pi^*)$. To see this, differentiating the LHS by $Z(\pi_t)$, we have:

$$-2Z(\pi_t) + \phi - c\phi - \pi_t a(\pi_t) \leq -2(\phi - c\phi - \varepsilon) + \phi - c\phi - \pi_t a(\pi_t) < 0.$$

So, if $Z_1(\pi_t) > Z_2(\pi_t)$ then $Z_1'(\pi_t) < Z_2'(\pi_t)$. So $Z_1(\pi_t) - Z_2(\pi_t)$ is increasing when π_t decreases, i.e.

$$\lim_{\pi_t \to 0} Z_1(\pi_t) - \lim_{\pi_t \to 0} Z_2(\pi_t) \geq Z_1 - Z_2 > 0,$$

contradicting $\lim_{\pi_t \to 0} Z_1(\pi_t) = \lim_{\pi_t \to 0} Z_2(\pi_t) = \phi - c\phi$.

A.4 Proof of Lemma 3.5

Let $V_a^*(\cdot)$ be the unique solution to (3.13). Re-write the ODE as:

$$\pi_t(1 - \pi_t) V_a^{*'}(\pi_t) = \frac{1}{ca^*(\pi_t)} \big(\phi - V_a^*(\pi_t)\big)\big(c\phi + \pi_t a^*(\pi_t) - V_a^*(\pi_t)\big), \tag{A.6}$$

This implies that $V_a^{*'}(\pi_t) > 0$ if and only if

$$V_a^*(\pi_t) < \pi_t a^*(\pi_t) + c\phi. \tag{A.7}$$

I show that $V_a^*(\pi_t)$ is strictly increasing in π_t for $\pi_t \in (0, \pi^\dagger)$. Suppose towards a contradiction that there exists $\pi < \pi^\dagger$ such that:

$$V_a^*(\pi) \geq \pi a^*(\pi) + c\phi,$$

then $V_a^{*'}(\pi) \leq 0$. Since $\pi_t a^*(\pi_t) + c\phi$ is a strictly increasing function of π_t and $V_a^*(\pi_t)$ is decreasing in π_t for all $\pi_t \in (0, \pi)$, so

$$V_a^*(\pi_t) - \Big(\pi_t a^*(\pi_t) + c\phi\Big)$$

strictly increases as π_t decreases when $\pi_t \leq \pi$. This implies that $\lim_{\pi_t \downarrow 0} V_a^*(\pi_t) > c\phi$, which contradicts the limiting value condition.

A.5 Proof of Lemma 3.6

According to Lemma 3.2, $\pi^1 > 0$. Suppose towards a contradiction that $\pi^1 \in (0, \min\{1 - b, \pi^\dagger\})$. Since $\pi^1 < \pi^\dagger$, Lemma 3.5 implies that $V_a(\pi^1) < \overline{V}_a$. So there exists $\varepsilon > 0$ such that $V_a(\pi^1 + \varepsilon_0) < \overline{V}_a$ for all $\varepsilon_0 \in (0, \varepsilon)$. Throughout the proof, I will be focusing on $\pi_t \in [\pi^1, \pi^1 + \varepsilon)$.

By definition of π^1, for every $\varepsilon > 0$, there exists $\pi_t \in (\pi^1, \pi^1 + \varepsilon)$ such that $a(\pi_t) \neq a^*(\pi_t)$. Since $a(\pi_t)$ is left-continuous, we only need to consider the following three cases. Common in all cases, there exists an open interval $\Pi \equiv (\pi_t - \varepsilon_0, \pi_t) \subset (\pi^1, \pi^1 + \varepsilon)$, such that:

- **Case 1:** $a(\pi) \in (\phi, 1)$ but $a(\pi) \neq a^*(\pi_t)$ for all $\pi \in \Pi$. By the choice of ε, the agent's continuation value is strictly below $\overline{V}_a$ for all $\pi \in \Pi$. So $\chi(\pi) \in (0,1)$ for all $\pi \in \Pi$. However, the intermediary's continuation value function requires that $a(\pi) = a^*(\pi)$ for all $\pi \in \Pi$, which is a contradiction.

- **Case 2:** $a(\pi) = 1$ for all $\pi \in \Pi$. Left-continuity of $a(\cdot)$ implies that $a(\pi_t) = 1$. Since $\chi(\pi) > 0$ for all $\pi \in \Pi \cup \{\pi_t\}$, we have $V_m(\pi) \leq b$.

  ODE (3.7) implies that if $V_m(\pi_t) \leq b$, then $V_m(\pi) < b$ for all $\pi \in \Pi$, and hence, $\chi(\pi) = 1$ by the intermediary's incentive constraint. Since the agent's continuation value is below $\overline{V}_a$ for all $\pi_t \leq \pi^1 + \varepsilon$, so $a(\pi) = 1$ whenever $\chi(\pi) = 1$. Let

  $$\pi^2 \equiv \inf \left\{ \pi \middle| \pi \in [\pi^1, \pi_t], \quad V_m(\pi') < b \text{ for all } \pi' > \pi \right\},$$

  then $\pi^2 = \pi^1$, implying that $V_m(\pi^1) < b$, contradicting the fact that $a(\pi^1) = a^*(\pi^1)$ and $\chi(\pi^1) \in (0,1)$.

- **Case 3:** $a(\pi) = \phi$ for all $\pi \in \Pi$. Left-continuity of $a(\cdot)$ implies that $a(\pi_t) = \phi$. Then, $\chi(\pi) < 1$ for all $\pi \in \Pi \cup \{\pi_t\}$, which requires that $V_m(\pi) \geq b$.

  ODE (3.7) implies that if $V_m(\pi_t) \geq b$, then $V_m(\pi) > b$ for all $\pi \in \Pi$, so $\chi(\pi) = 0$ by the intermediary's incentive constraint. This contradicts Lemma 3.1.

A.6 Proof of Lemma 3.7

I start with the following Lemma.

Lemma A.2. *If $\pi^2 \geq \overline{\pi}$, then $V_a(\pi_t) > \overline{V}_a$ for all $\pi_t > \pi^2$.*

Proof of Lemma A.2: Recall the agent's continuation value satisfies:

$$V_a(\pi_t) = \pi_t a_t - (a_t - \phi)c + \frac{\mu \chi_t a_t}{r}\Big(\phi - V_a(\pi_t) - \pi_t(1-\pi_t)V_a'(\pi_t)\Big). \tag{A.8}$$

By definition, if $\pi^2 \geq \overline{\pi}$, then $\chi(\pi_t) = 1$ and $V_a(\pi_t) \geq \overline{V}_a$ for all $\pi_t \in (\pi^\dagger, \pi^2]$. Let

$$\pi^* \equiv \inf \left\{ \pi_t \middle| \pi_t \in (\pi^2, 1) \text{ and } V_a(\pi_t) \leq \overline{V}_a \right\},$$

Suppose towards a contradiction that $\pi^* < 1$, then $V_a'(\pi^*) < 0$. The intermediary's incentives imply that $\chi(\pi^*) = 1$. Plugging this back to (A.8) and by the fact that $\pi^* \geq \pi^2 \geq \overline{\pi}$, we have $V_a'(\pi^*) \geq 0$, leading to a contradiction. □

Lemma A.3. *If $\pi^2 < \overline{\pi}$, then there exists $\pi_t \in (\pi^2, \overline{\pi})$ such that $V_a(\pi_t) \neq \overline{V}_a$.*

Proof of Lemma A.3: Suppose towards a contradiction that $V_a(\pi_t) = \overline{V}_a$ for all $\pi_t \in (\pi^2, \overline{\pi}) \equiv \Pi_0$. For any $\pi_t \in \Pi_0$ such that $a(\pi_t) \neq a^{**}(\pi_t)$, since $V_a(\pi_t) = \overline{V}_a$,

- If $\chi(\pi_t) = 1$, then $V_a'(\pi_t) \neq 0$, leading to a contradiction.

- If $\chi(\pi_t) < 1$, then $a(\pi_t) = \phi$, plugging this into (A.8) and using the assumption that $V_a'(\pi_t) = 0$,

  $$\overline{V}_a = \pi_t \phi + c\phi\chi_t < (1 - c - \frac{cr}{\mu\phi})\phi + c\phi = \phi - \frac{cr}{\mu} = \overline{V}_a,$$

  which leads to a contradiction.

This finishes the proof. $\square$

Back to Lemma 3.7, suppose towards a contradiction that $\pi^2 \in [\pi^\dagger, 1)$. If $\pi^2 > \overline{\pi}$, then according to Lemma A.2, $V_a(\pi_t) > \overline{V}_a$ for all $\pi_t > \pi^2$. In this case, only $a(\pi_t) = \phi = a^{**}(\pi_t)$ is incentive compatible, which leads to a contradiction.

When $\pi^2 < \overline{\pi}$. There exists $\pi > \pi^2$, such that $a(\pi) \neq a^{**}(\pi)$. Lemma A.3 implies that there exists $\pi_t \in (\pi^2, \overline{\pi})$ such that $V_a(\pi_t) \neq \overline{V}_a$. Two subcases are examined separately.

1. If $V_a(\pi_t) > \overline{V}_a$, then by the continuity of $V_a(\cdot)$ and Lemma 3.4, there exists $\pi_t' \in (\pi^2, \pi_t)$ such that
$$a(\pi_t') > \phi, \quad V_a(\pi') > \overline{V}_a.$$
This leads to a contradiction since when $V_a(\pi') > \overline{V}_a$, the only incentive compatible effort level is $a(\pi') = \phi$, regardless of $\chi(\pi')$.

2. If $V_a(\pi_t) < \overline{V}_a$, then there exists $\pi' \in (\pi^2, \pi_t)$ such that
$$a(\pi') < a^{**}(\pi'), \quad V_a(\pi') < \overline{V}_a.$$
Since $a(\cdot)$ is left-continuous and $V_a(\cdot)$ is continuous, there exists $\varepsilon_0 \in (0, \pi_t - \pi^2)$ such that
$$a(\pi'') < a^{**}(\pi''), \quad V_a(\pi'') < \overline{V}_a$$
for all $\pi'' \in (\pi' - \varepsilon_0, \pi') \equiv \Pi$. This also implies that $\chi(\pi'') < 1$ for all $\pi'' \in \Pi$. The intermediary's incentive constraint requires that $V_m(\pi' - \frac{\varepsilon_0}{2}) \geq b$.
From ODE (3.7), $V_m(\pi'') > b$ for all $\pi'' \in (\pi' - \varepsilon_0, \pi' - \frac{\varepsilon_0}{2})$. This implies that $\chi(\pi'') = 0$, contradicting the conclusion of Lemma 3.1.

A.7 Proof of Lemma 3.8

I start from the following Lemma:

Lemma A.4. *If $\pi_0 < \phi$, then for any $t \in \mathbb{R}_+$, $\boldsymbol{\chi} \in \mathbf{X}$, $\mathbf{a} \in \mathbf{A}$ and $\hat{\mathbf{a}} \in \mathbf{A}$:*
$$\mathbb{E}^{\hat{\mathbf{a}}, \boldsymbol{\chi}, \pi_0}[\pi_{t+dt} | \theta = 1, \mathbf{a}] > \pi_t. \tag{A.9}$$

Proof of Lemma A.4: Conditional on $\theta = 1$, the probability breakthrough at $[t, t+dt]$ is at least $\mu\chi_t\phi dt$, after which $\pi_t = 1$. With complementary probability, π_t degrades to $\pi_t - \mu\chi_t\hat{a}_t\pi_t(1-\pi_t)dt$. The expected belief at $t + dt$ exceeds π_t if:
$$\mu\chi_t\phi dt + (1 - \mu\chi_t\phi dt)(\pi_t - \mu\chi_t\hat{a}_t\pi_t(1-\pi_t)dt) > \pi_t \tag{A.10}$$
Ignoring higher order terms, we get:
$$\phi > \hat{a}_t\pi_t, \tag{A.11}$$
(A.11) implies that $\pi_0 < \phi$ is sufficient for (A.9). $\square$

Back to the proof of Lemma 3.8, for any given $\boldsymbol{\chi}$ and a *correct* market belief $\hat{\mathbf{a}}$, the agent's continuation value is:
$$V_a^{\hat{\mathbf{a}}, \boldsymbol{\chi}}(\pi_t) = \pi_t\hat{a}(\pi_t) - c\Big(a(\pi_t) - \phi\Big) + \frac{\mu\chi(\pi_t)a(\pi_t)}{r}\Big(\phi - V_a^{\hat{\mathbf{a}}, \boldsymbol{\chi}}(\pi_t)\Big) - \frac{\mu\chi(\pi_t)\hat{a}(\pi_t)}{r}\pi_t(1-\pi_t)V_a^{\hat{\mathbf{a}}, \boldsymbol{\chi}\prime}(\pi_t). \tag{A.12}$$

Let $\Gamma^{\hat{\mathbf{a}},\boldsymbol{\chi}}(\pi_t) \equiv V_a^{\hat{\mathbf{a}},\boldsymbol{\chi}}(\pi_t) - \pi_t\phi$, we have:

$$\begin{aligned}\Gamma^{\hat{\mathbf{a}},\boldsymbol{\chi}}(\pi_t) = \pi_t\Big(\hat{a}(\pi_t) - \phi\Big) - c\Big(a(\pi_t) - \phi\Big) + \frac{\mu\chi(\pi_t)a(\pi_t)}{r}\Big(\phi(1-\pi_t) - \Gamma^{\hat{\mathbf{a}},\boldsymbol{\chi}}(\pi_t)\Big) \\ - \frac{\mu\chi(\pi_t)\hat{a}(\pi_t)}{r}\pi_t(1-\pi_t)\Gamma^{\hat{\mathbf{a}},\boldsymbol{\chi}\prime}(\pi_t) - \frac{\mu\phi\chi(\pi_t)\hat{a}(\pi_t)}{r}\pi_t(1-\pi_t). \qquad \text{(A.13)}\end{aligned}$$

In what follows, I write $\Gamma(\cdot)$ instead of $\Gamma^{\hat{\mathbf{a}},\boldsymbol{\chi}}(\cdot)$ for notation simplicity. Let

$$\pi^* \equiv \sup\Big\{\pi_t \leq 1 \Big| V_a^{\hat{\mathbf{a}},\boldsymbol{\chi}}(\pi) \leq \phi\pi \text{ for all } \pi \leq \pi_t\Big\}.$$

Lemma A.4 suggests that when $\pi_0 < \phi$, the agent can guarantee himself payoff $\pi_0\phi$ by playing $a_t = \phi$ forever, so $\pi^* \geq \phi > 0$.

I show that $\pi^* = 1$. Suppose towards a contradiction that $\pi^* < 1$. Since $\Gamma(\pi_t)$ is continuous, we have:

$$\Gamma(\pi^*) = 0, \quad \Gamma'(\pi^*) < 0.$$

From (A.13), we have:

$$0 > \pi^*\Big(\hat{a}(\pi^*) - \phi\Big) - c\Big(a(\pi^*) - \phi\Big) + \frac{\mu\chi(\pi^*)a(\pi^*)}{r}\phi(1-\pi^*) - \frac{\mu\phi\chi(\pi^*)\hat{a}(\pi^*)}{r}\pi^*(1-\pi^*). \qquad \text{(A.14)}$$

I consider three cases, depending on the market's believed effort $\hat{a}$ as well as the strength of reputation concerns, $\frac{\mu\chi(\pi^*)}{r}\phi(1-\pi^*) - c$. I will obtain a contradiction in each case based on (A.14).

Case 1: If $\hat{a}(\pi^*) = \phi$, then $a(\pi^*) = \phi$ must be a best reply since the market's belief is correct, we have:

$$\begin{aligned} 0 &> \frac{\mu\chi(\pi^*)\phi}{r}\phi(1-\pi^*) - \frac{\mu\phi\chi(\pi^*)}{r}\phi\pi^*(1-\pi^*) \\ &= \frac{\mu\chi(\pi^*)\phi^2}{r}(1-\pi^*)^2 \geq 0, \end{aligned}$$

which is a contradiction.

Case 2: If $\hat{a}(\pi^*) \in [\phi, 1]$ and $\frac{\mu\chi(\pi^*)}{r}\phi(1-\pi^*) = c$, then the agent is indifferent between all effort levels, so

$$\begin{aligned} 0 &> \pi^*\Big(\hat{a}(\pi^*) - \phi\Big) + c\phi - \underbrace{\frac{\mu\phi\chi(\pi^*)\hat{a}(\pi^*)}{r}\pi^*(1-\pi^*)}_{=c\pi^*\hat{a}(\pi^*)}. \\ &= (1-c)\pi^*\hat{a}(\pi^*) + \phi(c - \pi^*) \\ &\geq (1-c)\pi^*\phi + \phi(c-\pi^*) = \phi c(1-\pi^*) > 0. \end{aligned}$$

which is a contradiction.

Case 3: If $\hat{a}(\pi^*) = 1$ and $\frac{\mu\chi(\pi^*)}{r}\phi(1-\pi^*) > c$, then $a(\pi^*) = 1$ is the agent's best response. Then

$$\begin{aligned} 0 &> \pi^*\Big(1-\phi\Big) - c\Big(1-\phi\Big) + \frac{\mu\chi(\pi^*)}{r}\phi(1-\pi^*) - \frac{\mu\phi\chi(\pi^*)}{r}\pi^*(1-\pi^*). \\ &= (\pi^*-c)(1-\phi) + \frac{\mu\phi\chi(\pi^*)}{r}(1-\pi^*)^2 \\ &> (\pi^*-c)(1-\phi) + c(1-\pi^*) = \pi^*(1-c) + \phi(c-\pi^*). \end{aligned}$$

If $\pi^* \le c$, the last expression is strictly positive, which leads to a contradiction. If $\pi^* > c$ and $\pi^* \ge \phi$,

$$\pi^*(1-c) + \phi(c-\pi^*) \ge \pi^*(1-c) + \pi^*(c-\pi^*) = \pi^*(1-\pi^*) > 0,$$

which leads to a contradiction. If $\pi^* > c$ but $\pi^* < \phi$,

$$\pi^*(1-c) + \phi(c-\pi^*) \ge \pi^*(1-c) + (c-\pi^*) = c(1-\pi^*) > 0,$$

which also leads to a contradiction.

B Exogenous Information Benchmark

In this Appendix, I show Proposition 1 and Proposition 3.

B.1 Proof of Proposition 1

I start with computing the agent's continuation value when $\pi_t \to 0$ if $\chi_t = 1$ for all t. Applying integral expression (A.3) and plugging in $\chi_t = 1$, we have:

$$\hat{V}_a(t) - c\phi = X_t + W_t.$$

For any $\varepsilon > 0$, there exists $T \in \mathbb{R}_+$ such that $W_t \in (0, \varepsilon)$ for all $t > T$.

Suppose towards a contradiction that $\hat{V}_a(\infty) \le c\phi$, then $X_t \to 0$ and there exists $\varepsilon > 0$, such that:

$$\frac{\mu}{r}(\phi - \hat{V}_a(t)) > \frac{\mu}{r}(\phi - c\phi) > c + \varepsilon,$$

for t large enough. Hence, $Z_t \ge \varepsilon$, i.e. $X_t \ge \varepsilon$ for t large enough. But then,

$$\lim_{t\to\infty} \hat{V}_a(t) = \lim_{t\to\infty} W_t + \lim_{t\to\infty} X_t + c\phi \ge c\phi + \varepsilon > c\phi,$$

leading to a contradiction.

So $\hat{V}_a(\infty) > c\phi$, implying that $\lim_{t\to\infty} X_t > 0$, i.e. there exists $T \in \mathbb{R}_+$ such that $Z_t > 0$ for all $t > T$. Thus, $\hat{V}_a(\infty)$ satisfies:

$$\hat{V}_a(\infty) = c\phi + \lim_{t\to\infty} X_t = c\phi + r\int_0^\infty \Big(\frac{\mu}{r}(\phi - \hat{V}_a(\infty)) - c\Big)dt, \tag{B.1}$$

which gives:

$$\hat{V}_a(\infty) = V_a(0) = \frac{\mu\phi}{\mu + r} - \frac{rc(1-\phi)}{\mu + r}. \tag{B.2}$$

Next, I construct an MPE of the exogenous information game and apply the result of Faingold and Sannikov (2011) to establish uniqueness.[44] Consider the following limiting value problem:

$$V_a(\pi_t) = \pi_t - c(1-\phi) + \frac{\mu}{r}\Big(\phi - V_a(\pi_t) - \pi_t(1-\pi_t)V_a'(\pi_t)\Big),$$

with (B.2) the limiting value condition. I will show that this problem admits a unique solution in section A of the Online Appendix. Let $V_a^{**}(\pi_t)$ be the solution. Let $\pi^{\ddagger}$ be defined by:

$$\pi^{\ddagger} \equiv \inf\Big\{\pi_t > 0 \Big| V_a^{**}(\pi_t) = \overline{V}_a\Big\}.$$

From Lemma 3.4, $\pi^{\ddagger} > \underline{\pi}$. From Lemma 3.8, $\pi^{\ddagger} < 1 - \frac{cr}{\mu\phi}$. Let

$$a(\pi_t) = \begin{cases} 1 & \text{when } \pi_t \geq \pi^{\ddagger} \\ a^{**}(\pi_t) & \text{when } \pi_t < \pi^{\ddagger}. \end{cases} \tag{B.3}$$

I claim that this strategy and its induced market belief system induce an MPE. First, $V_a^{**}(\pi_t)$ is the agent's continuation value when $\pi_t \leq \pi^{\ddagger}$. By definition of $\pi^{\ddagger}$, $V_a^{**}(\pi_t) \leq \overline{V}_a$ for all $\pi_t \leq \pi^{\ddagger}$, implying that $a(\pi_t) = 1$ is incentive compatible for the agent. Second, since $V_a(\pi^{\ddagger}) = \overline{V}_a$, so

- If $\pi^{\ddagger} < \overline{\pi}$, then since $a(\pi_t) = a^{**}(\pi_t)$, we have $V_a(\pi_t) = \overline{V}_a$ for all $\pi_t \in [\pi^{\ddagger}, \overline{\pi}]$ (Lemma 3.4), implying that every effort level is optimal. If $\pi_t > \overline{\pi}$, then $V_a(\pi_t) > \overline{V}_a$, implying that $a(\pi_t) = \phi$ is optimal for the agent.
- If $\pi^{\ddagger} \geq \overline{\pi}$, then $V_a(\pi_t) > \overline{V}_a$ for all $\pi_t > \pi^{\ddagger}$, i.e. $a(\pi_t) = \phi$ is optimal for the agent.

B.2 Proof of Proposition 3

Let $V_1(\cdot)$ be the unique solution to the following limiting value problem:

$$V_a(\pi) = \pi a^*(\pi) + c\phi - \frac{c}{\phi - V_a(\pi)}\pi(1-\pi)V_a'(\pi),$$

with $\lim_{\pi\to 0} V_a(\pi) = c\phi$. Let $V_2(\cdot)$ be the unique solution to the following limiting value problem:

$$V_a(\pi) = \pi - c(1-\phi) + \frac{\mu}{r}\Big(\phi - V_a(\pi) - \pi(1-\pi)V_a'(\pi)\Big),$$

with $\lim_{\pi\to 0} V_a(\pi) = \frac{r}{\mu+r}\overline{V}_a + \frac{\mu}{\mu+r}c\phi$.

Recall that $\pi^{\dagger}$ and $\pi^{\ddagger}$ are defined by $V_1(\pi^{\dagger}) = \overline{V}_a$ and $V_2(\pi^{\ddagger}) = \overline{V}_a$ respectively. The following Lemma puts upper bounds on $V_1(\cdot)$ and $V_2(\cdot)$.

Lemma B.1. *For all $\pi_t \in (0, \pi^{\dagger}]$,*

$$V_1(\pi_t) \leq \pi_t a^*(\pi_t) + c\phi.$$

For all $\pi_t \in (0, \pi^{\ddagger}]$,

$$V_2(\pi_t) \leq \frac{r}{\mu+r}\pi_t + \frac{r}{\mu+r}\overline{V}_a + \frac{\mu}{\mu+r}c\phi.$$

[44] The exogenous information benchmark fits into the definition of Poisson good news model in Faingold and Sannikov (2011), subsection 9.1 page 823, and the uniqueness result is reported in Theorem 11 of their paper.

This Lemma will be shown in the Online Appendix.[45] Let

$$\pi^* \equiv \frac{2\mu + r}{2\mu + 2r}\phi, \tag{B.4}$$

which is less than $\frac{1}{2}$ when ϕ is small enough ($\phi < \frac{1}{2}$). Since $V_1(\pi^*) < c\phi + \pi^* a(\pi^*) < c\phi + \pi^*$, we have $V_1(\pi^*) < \overline{V}_a$. I will be focusing on the case in which $b = 1 - \pi^*$. The claim can be extended to an open neighborhood of b by a continuity argument. The key step is the following Lemma:

Lemma B.2. *When $\phi < \frac{1}{2}$, for every $\underline{r} > 0$, there exists $\overline{c} \in (0,1)$ and $\underline{\mu} > 0$ satisfying*

$$\phi - \underline{c}\phi - \frac{r\overline{c}}{\underline{\mu}} > 0$$

such that for every $c < \overline{c}$, $\mu > \underline{\mu}$ and $r \in \left(\underline{r}, \frac{\mu\phi(1-\overline{c})}{\overline{c}}\right)$, we have:

$$V_2(\pi^*) < V_1(\pi^*). \tag{B.5}$$

Proof of Lemma B.2: The proof is divided into two parts. In Part I, I show that $V_1(\pi^*)$ is arbitrarily close to π^* when $c \to 0$ and r is sufficiently large. In Part II, I show that when μ and r are both large, $V_2(\pi^*)$ is bounded away from π^* even when $c \to 0$.

Part I: Lemma B.1 implies that $V_1(\pi^*) \le \pi^* + c\phi$, so

$$\overline{V}_a - V_1(\pi^*) \ge \phi - c\phi - \frac{cr}{\mu} - \pi^* = \frac{r\phi}{2(\mu + r)} - (c\phi + \frac{cr}{\mu}).$$

Since

$$\frac{\mu\chi(\pi^*)}{r}\Big(\phi - V_1(\pi^*)\Big) = c,$$

we have:

$$\chi(\pi^*) \le \frac{cr}{\mu}\frac{1}{\phi\frac{r}{2\mu + 2r} - c(\phi + \frac{r}{\mu})}.$$

Notice that the agent's flow payoff is bounded above by 1 and below by -1. For any $r > 0$ and $\varepsilon_0 > 0$, there exists $T > 0$ such that

$$r\int_0^T e^{-rt}dt > 1 - \varepsilon_0.$$

Let π_t be defined via the following ODE

$$\dot{\pi}_t = -\mu\pi_t(1 - \pi_t)\chi(\pi_t), \tag{B.6}$$

with initial value $\pi_0 = \pi^*$. Then for every $\varepsilon_1 > 0$, there exists c small enough and μ large enough such that

$$\pi^* - \pi_T < \varepsilon_1.$$

Also, by definition of π^*, which is less than ϕ, the agent's flow payoff at time t is at least π_t. Now, we can compute a lower bound on $V_1(\pi^*)$:

$$V_1(\pi^*) \ge r\int_0^T e^{-rt}\min\Big\{\pi^* - \varepsilon_1, \phi\Big\}dt + r\int_T^\infty e^{-rt}(-1)dt \ge (\pi^* - \varepsilon_1)(1 - \varepsilon_0) - \varepsilon_0,$$

which converges to π^* as both ε_0 and ε_1 go to 0.

[45] The first inequality is a Corollary of Lemma A.3 in the Online Appendix. The second inequality is a Corollary of Lemma A.4 in the Online Appendix.

Part II: I will show that $V_2(\pi^*)$ is bounded below by π^*, with difference bounded away from 0 even when $c \to 0$. This can imply that $V_2(\pi^*) < V_1(\pi^*)$. I establish an upper bound on t^*, defined via $\pi_{t^*} = \frac{\phi}{3}$, where π_t is the solution to initial value problem (B.6).

$$t^* \leq \frac{\frac{2\mu+r}{2\mu+2r}\phi - \frac{\phi}{3}}{\frac{\phi}{3}(1-\frac{\phi}{3})\mu} = \frac{1}{2(1-\frac{\phi}{3})}\frac{4\mu+r}{\mu(\mu+r)}.$$

Intuitively, t^* is the time it takes for market to belief to go from π^* to $\frac{\phi}{3}$. Let $C_1 \equiv \frac{1}{2(1-\frac{\phi}{3})}$. We can then obtain an upper bound on $V_2(\pi^*)$:

$$\begin{aligned} V_2(\pi^*) &\leq r\int_0^{C_1\frac{4\mu+r}{\mu(\mu+r)}} e^{-rt}dt + V_2(\frac{\phi}{3})\Big(1 - r\int_0^{C_1\frac{4\mu+r}{\mu(\mu+r)}} e^{-rt}dt\Big) \\ &\leq \Big(1-e^{-rC_1\frac{4\mu+r}{\mu(\mu+r)}}\Big) + e^{-rC_1\frac{4\mu+r}{\mu(\mu+r)}}\Big(\underbrace{\frac{\phi}{3}\frac{r}{\mu+r} + \frac{\mu\phi}{\mu+r}}_{=\pi^* - \frac{\phi}{6}\frac{r}{\mu+r}} - \frac{rc(1-\phi)}{\mu+r}\Big). \end{aligned}$$

For every $\varepsilon > 0$, there exists $\underline{\mu} > 0$ such that for all $\mu > \underline{\mu}$,

$$\Big(1-e^{-rC_1\frac{4\mu+r}{\mu(\mu+r)}}\Big) + (\frac{\phi}{2} + c\phi)e^{-rC_1\frac{4\mu+r}{\mu(\mu+r)}} \leq \pi^* - \frac{\phi}{6}\frac{r}{r+\mu} + \varepsilon.$$

□

Lemma B.2 has found a belief, π^*, such that $V_2(\pi^*) < V_1(\pi^*) < \overline{V}_a$. What remains to be shown is that $V_1(\pi_t) > V_2(\pi_t)$ for all $\pi_t \in [\pi^*, \min\{\pi^\dagger, \pi^\ddagger\}]$. Since $1-b = \pi^*$, so when $\pi_t > \pi^*$, $(a_t, \chi_t) = (1,1)$ in both scenarios, implying that $V_1(\cdot)$ reaches $\overline{V}_a$ at a smaller belief than $V_2(\cdot)$.[46]

C Semi-Markov Equilibrium

C.1 Characterizing π^* and π^{**}

In this subsection, I characterize the two belief thresholds, π^* and π^{**}, as functions of the cutoff $\pi^\S$, which we write as $\pi^*(\pi^\S)$ and $\pi^{**}(\pi^\S)$, respectively. I characterize the agent's continuation value in each phase, as well as the thresholds π^* and π^{**}, as functions of $\pi^\S$. Let $V_a^*(\cdot)$ be the unique solution to limiting value problem (3.13), which is the agent's continuation value in Phase I.

Recall that $\pi^\dagger$ is defined in (3.14) by $V_a^*(\pi^\dagger) = \overline{V}_a$. For any $\pi^\S \in [0, \min\{1-b, \pi^\dagger\}]$, Lemma 3.5 implies that $V_a^*(\pi^\S) < \overline{V}_a$. Let $V_a^1(\cdot|\pi^\S)$ be the unique solution to the following initial value problem:

$$V_a(\pi_t) = \pi_t - c(1-\phi) + \frac{\mu}{r}\Big(\phi - V_a(\pi_t) - \pi_t(1-\pi_t)V_a'(\pi_t)\Big), \tag{C.1}$$

with $V_a(\pi^\S) = V_a^*(\pi^\S)$. By definition, $V_a^1(\pi_t|\pi^\S)$ is the agent's continuation value in Phase II when $\pi_t \leq \pi^*(\pi^\S)$, with $\pi^*(\pi^\S)$ defined as:

$$\pi^*(\pi^\S) \equiv \sup\Big\{\pi_t \Big| \pi_t \geq \pi^\S,\ V_a^1(\pi|\pi^\S) < \overline{V}_a \text{ for all } \pi < \pi_t\Big\}. \tag{C.2}$$

[46]This perverse incentive also exists when the intermediary is the agent's direct supervisor, i.e. her flow payoff is θa_t instead of $\theta a_t - w_t$. The proof of this extension is available upon request.

If this set is empty, then $\pi^*(\pi^\S) = \pi^\S$. By definition,

$$\pi^*(0) = \pi^\ddagger \text{ and } \pi^*\Big(\min\{1-b, \pi^\dagger\}\Big) = \pi^\dagger.$$

Let $a^1(\pi_t|\pi^\S)$ be the agent's on-path effort given that punishment phase starts at $\pi^\S$. Let $Y^1(\pi_t|\pi^\S) \equiv \mu a^1(\pi_t|\pi^\S)$ be market's on-path learning rate. Let $\hat{V}_a^0(\cdot|\pi^\S)$ be the (unique) solution to the following initial value problem:

$$V_a(\pi_t) = \pi_t a^1(\pi_t|\pi^\S) + c\phi - \frac{Y^1(\pi_t|\pi^\S)}{r}\pi_t(1-\pi_t)V_a'(\pi_t). \tag{C.3}$$

with $V_a(\pi^\S) = V_a^*(\pi^\S)$. Define $\hat{\pi}^{**}(\pi^\S)$ as:

$$\hat{\pi}^{**}(\pi^\S) \equiv \sup\Big\{\pi_t \Big| \pi_t \geq \pi^\S,\ \hat{V}_a^0(\pi|\pi^\S) < \overline{V}_a \text{ for all } \pi < \pi_t\Big\}. \tag{C.4}$$

If this set is empty, then $\hat{\pi}^{**}(\pi^\S) = \pi^\S$. Suppose $\hat{\pi}^{**}(\pi^\S) \leq 1-b$, then

$$\pi^{**}(\pi^\S) \equiv \hat{\pi}^{**}(\pi^\S). \tag{C.5}$$

Otherwise, record the value $\hat{V}_a^0(1-b|\pi^\S)$. Let $\tilde{V}_a^0(\cdot|\pi^\S)$ be the (unique) solution to the following initial value problem:

$$V_a(\pi_t) = \pi_t a^1(\pi_t|\pi^\S) + c\phi + \frac{\mu}{r}\Big(\phi - V_a(\pi_t)\Big) - \frac{Y^1(\pi_t|\pi^\S)}{r}\pi_t(1-\pi_t)V_a'(\pi_t). \tag{C.6}$$

with $V_a(1-b) = \hat{V}_a^0(1-b|\pi^\S)$. Define $\tilde{\pi}^{**}(\pi^\S)$ as:

$$\tilde{\pi}^{**}(\pi^\S) \equiv \sup\Big\{\pi_t \Big| \pi_t \geq 1-b,\ \tilde{V}_a^0(\pi|\pi^\S) < \overline{V}_a \text{ for all } \pi < \pi_t\Big\}. \tag{C.7}$$

Then we have:

$$\pi^{**}(\pi^\S) \equiv \tilde{\pi}^{**}(\pi^\S). \tag{C.8}$$

C.2 Proof of Proposition 4

First, the proof of Lemma 3.1 directly carries over to the Semi-Markov case which implies that $\pi_t(1-\pi_t) \to 0$ in probability. I use $\{a(\pi_t), \chi(\pi_t)\}$ to represent players' on-path strategies. The next Lemma is the counterpart of Lemma 3.2 which characterizes the agent's continuation value in a generic SME when $\pi_t \to 0$.

Lemma C.1. *When players' patience level is high, either one of the following two statements is true:*[47]

- $\lim_{\pi_t \to 0} V_a(\pi_t) = c\phi$ *and there exists* $\nu > 0$ *such that for all* $\pi_t \leq \nu$, $a(\pi_t) = a^*(\pi_t)$ *and* $\chi(\pi_t) < 1$.
- $\lim_{\pi_t \to 0} V_a(\pi_t) = \frac{\mu\phi - rc(1-\phi)}{\mu + r}$ *and there exists* $\nu > 0$ *such that for all* $\pi_t \leq \nu$, $a(\pi_t) = 1$ *and* $\chi(\pi_t) = 1$.

Comparing with Lemma 3.2, SME admits another possibility (statement 2). This is because $(a_t, \chi_t) = (1, 1)$ can be sustained by low effort low disclosure rate off the equilibrium path.

[47]Recall that $c\phi$ is the agent's limiting continuation value in the unique MPE and $\frac{\mu\phi - rc(1-\phi)}{\mu + r}$ is his limiting continuation value under exogenous information.

Proof of Lemma C.1: Let $V_a(\pi_t), V_m(\pi_t), a(\pi_t)$ and $\chi(\pi_t)$ be the agent's and the intermediary's continuation value, effort and disclosure rate on the equilibrium path. Let $\hat{V}_a(t) \equiv V_a(\pi_t)$, $\hat{V}_m(t) = V_m(\pi_t)$. Define X_t, Z_t and W_t as in the proof of Lemma 3.2, we have:

$$\hat{V}_a(t) \quad = \quad c\phi + X_t + W_t. \tag{C.9}$$

Since the only relevant off-path is the intermediary withholding information when $\chi(\pi_t) = 1$, then the proofs in Part I, II and IV of Lemma 3.2 directly go through. I modify the argument in Part III.

Modified Part III: Suppose $V_a(0) \in (c\phi, \overline{V}_a)$.[48] Then there exists $T \in \mathbb{R}_+$ such that $\hat{V}_a(t) < \overline{V}_a$ and $X_t > 0$ for all $t > T$. So, there exists $t^* > T$ such that:

$$a_{t^*} = 1 \text{ and } \frac{\mu\chi_{t^*}}{r}\left(\phi - \hat{V}_a(t^*)\right) > c.$$

This also suggests that $\chi_t > 0$, which implies $V_m(\pi_t) \leq b$. Two cases are considered

- Suppose $\chi_t \neq 1$, then the Semi-Markov restriction implies that continuation play should not depend on whether a breakthrough has arrived or not, conditional on no disclosure. Admissibility requires the existence of $\varepsilon_0 > 0$ such that

  $$a(\pi_{t+\varepsilon_1}) \in (a^*(\pi_{t+\varepsilon_1}), 1) \text{ for all } \varepsilon_1 \in (0, \varepsilon_0).$$

  ODE (3.7) implies that $V_m(\pi_{t+\varepsilon_1}) < b$ for all $\varepsilon_1 \in (0, \varepsilon_0)$, which further implies that $\chi_{t+\varepsilon_1} = 1$ for all $\varepsilon_1 \in (0, \varepsilon_0)$. Since $\hat{V}_a(t) < \overline{V}_a$ for all $t > T$, so $a_t = 1$ and $\chi_t = 1$. But then $1 = V_m(\pi_t) > b$, contradicting $V_m(\pi_t) \leq b$. This contradiction applies as long as for every $T \in \mathbb{R}^+$, there exists $t > T$ such that

  $$Z_t > 0 \text{ and } \chi_t < 1.$$

- Next I consider the case in which $\chi_t = 1$ for all $t > T$ satisfying $Z_t > 0$.

  - If there exists $T \in \mathbb{R}^+$, such that $a_t = 1$ for all $t \geq T$, then we have:

    $$V_a(0) = \frac{\mu\phi - rc(1-\phi)}{\mu + r}.$$

  - If for every $T \in \mathbb{R}^+$, there exists $t > T$ such that $a(\pi_t) < 1$. Since $\hat{V}_a(t) < \overline{V}_a$, it is required that $\chi_t < 1$ and $V_m(\pi_t) = b$. Since $X_t > 0$, for every such t, there exists $t' > t$ such that $a(\pi_{t'}) = 1$. Admissibility requires that for every t, there exists $\varepsilon > 0$ such that

    $$a(\pi_{t+\varepsilon_0}) < 1 \text{ for all } \varepsilon_0 \in (0, \varepsilon).$$

    This requires that $\chi(\pi_{t+\varepsilon_0}) \in (0, 1)$, i.e. $V_m(\pi_{t+\varepsilon_0}) = b$. ODE (3.7) then requires that $a(\pi_{t+\varepsilon_0}) = a^*(\pi_{t+\varepsilon_0})$.
    This suggests that $a(\pi_t) \geq a^*(\pi_t)$ for all $t \geq T$. Also, for every t such that $a(\pi_t) = a^*(\pi_t)$, there exists $t' > t$ with $a(\pi_{t'}) = 1 > a^*(\pi_{t'})$. But then $V_m(\pi_t) > b$, implying that $\chi_t = 0$, contradicting the conclusion that $\chi_t \neq 0$.

□

[48] This case also treats the situation in which $V_a(0) = \overline{V}_a$ but $\hat{V}_a(t) < \overline{V}_a$ for all t large enough.

Next, I characterize players' *on-path* behaviour when π_t is bounded away from 0. The next two Lemmas examine cases when $V_a(\pi_t) < \overline{V}_a$.

Lemma C.2. *For any π_t such that $V_a(\pi_t) < \overline{V}_a$, $a(\pi_t) \in \{a^*(\pi_t), 1\}$.*

Proof of Lemma C.2: Suppose towards a contradiction that there exists such π_t. Then let

$$\pi^* \equiv \inf\Big\{\pi\Big|V_a(\pi) < \overline{V}_a, \quad a(\pi_t) \notin \{a^*(\pi_t), 1\}\Big\}. \tag{C.10}$$

By left-continuity of $a(\cdot)$, there exists $\varepsilon > 0$ such that $a(\pi^* + \varepsilon_0) \notin \{a^*(\pi^* + \varepsilon_0), 1\}$ for all $\varepsilon_0 \in (0, \varepsilon)$. Since $V_a(\pi_t) < \overline{V}_a$, so $\chi(\pi_t) < 1$ for all $\pi_t \in (\pi^*, \pi^* + \varepsilon)$. Since $\chi(\pi_t) \neq 0$, so $\chi(\pi_t) \in (0, 1)$, implying that $V_m(\pi_t) = b$ for all $\pi_t \in (\pi^*, \pi^* + \varepsilon)$. But the intermediary's continuation value satisfies:

$$V_m(\pi_t) = a(\pi_t)(1 - \pi_t) + \frac{\mu\chi(\pi_t)a(\pi_t)}{r}\Big\{b - V_m(\pi_t) - \pi_t(1 - \pi_t)V_m'(\pi_t)\Big\}.$$

Then $V_m(\pi_t) = b$ and $V_m'(\pi_t) = 0$ imply that $a(\pi_t) = a^*(\pi_t)$, leading to a contradiction. □

Lemma C.3. *If there exists π^* such that $a(\pi^*) = \chi(\pi^*) = 1$, then for any $\pi_t > \pi^*$ with $V_a(\pi_t) < \overline{V}_a$, we have $a(\pi_t) = \chi(\pi_t) = 1$.*

Proof of Lemma C.3: Suppose towards a contradiction that there exists such π_t, then $a(\pi_t) < 1$ while $V_a(\pi_t) < \overline{V}_a$ implies that $\chi(\pi_t) < 1$. Let

$$\pi^{**} \equiv \inf\Big\{\pi\Big|\pi > \pi^*, \quad V_a(\pi) < \overline{V}_a, \quad \chi(\pi) < 1, \quad a(\pi_t) < 1\Big\}. \tag{C.11}$$

By the left-continuity of $a(\cdot)$ and $\chi(\cdot)$, there exists $\varepsilon > 0$ such that $a(\pi^{**} + \varepsilon_0) < 1$ and $\chi(\pi^{**} + \varepsilon_0) < 1$ for all $\varepsilon_0 \in (0, \varepsilon)$. Then since withholding information happens on path before π_t reaches π^{**}, so according to the Semi-Markov restriction as well as Lemma C.2, $V_m(\pi^{**}) > b$. But then, $\chi(\pi^{**}) = 0$, which is a contradiction. □

Putting together Lemma C.1, Lemma C.2 and Lemma C.3, for all π_t such that $V_a(\pi_t) < \overline{V}_a$, there exists a cut-off belief $\pi^\S \in [0, \min\{1 - b, \pi^\dagger\}$, such that

$$a(\pi_t) = \begin{cases} a^*(\pi_t) & \text{when } \pi_t \leq \pi^\S \\ 1 & \text{when } \pi_t > \pi^\S \end{cases}$$

$$\chi(\pi_t) = \begin{cases} \dfrac{cr}{\mu(\phi - V_a(\pi_t))} & \text{when } \pi_t \leq \pi^\S \\ 1 & \text{when } \pi_t > \pi^\S \end{cases}$$

The limiting value problem in (3.13) and the initial value problem in (C.1) pin down the agent's on-path continuation value. Once it reaches $\overline{V}_a$ at $\pi^*(\pi^\S)$, then the proofs for uniqueness of on path effort and disclosure rate are exactly the same as in the proof of Lemma 3.7.

The last step verifies the players' incentives in Phase III. The agent's incentive to choose $a(\pi_t) = a^*(\pi_t)$ when $\pi_t \in (\pi^\S, \pi^{**}]$ and the intermediary's incentive when $\pi_t \leq \min\{1 - b, \pi^{**}\}$ is straightforward. To verify the agent's incentive to choose $a(\pi_t) = \phi$ when $\pi_t > \pi^{**}$, we compute his value function:

$$V_a(\pi_t) = \pi_t a(\pi_t) + c\phi + \hat{a}(\pi_t)\Big(\frac{\mu}{r}(\phi - V_a(\pi_t)) - c\Big) - \frac{\mu a(\pi_t)}{r}\pi_t(1 - \pi_t)V_a'(\pi_t), \tag{C.12}$$

where $a(\pi_t)$ and $\hat{a}(\pi_t)$ are on-path and off-path effort, respectively. When $\pi_t = \pi^{**}$, we have $V_a(\pi_t) = \overline{V}_a$, which gives:

$$\overline{V}_a = \pi_t a(\pi_t) + c\phi - \frac{\mu a(\pi_t)}{r}\pi_t(1 - \pi_t)V_a'(\pi_t),$$

Since $a(\pi_t) \geq a^{**}(\pi_t)$ for all π_t, this implies that $V_a'(\pi_t) > 0$ and the agent has an incentive to shirk at belief higher than π^{**}.