A Dynamic Model of Censorship

Yiman Sun*

January 12, 2019

Abstract

We model censorship as a dynamic game between an agent and an evaluator. Two types of public news are informative about the agent's ability – a conclusive good news process and a bad news process. However, the agent can hide bad news from the evaluator, at some cost, and will do so if and only if this secures her a significant increase in tenure. Thus, the evaluator faces a bandit problem with an endogenous news process. When bad news is conclusive, the agent always censors when the public belief is sufficiently high, but below a threshold, she either stops censoring or only censors with some probability, depending on the information structure. The possibility of censorship hurts the evaluator and the good agent, and it may also hurt the bad agent. However, when bad news is inconclusive, we show that the good agent censors bad news more aggressively than the bad agent does. This improves the quality of information, and may benefit all players – the evaluator, the bad agent and the good agent.

Keywords: Censorship, Information Manipulation, Learning, Dynamic Games

JEL Codes: C73, D82, D83

*Department of Economics, University of Texas at Austin. Email: yiman@utexas.edu
I am very grateful to my advisors, V. Bhaskar and Caroline Thomas, for their constant support and invaluable guidance. I would like to thank Daniel Ackerberg, Svetlana Boyarchenko, Yi Chen, Cary Deck, Tommaso Denti, Laura Doval, Ignacio Esponda, Andrew Glover, Yingni Guo, Ayça Kaya, Frédéric Koessler, Stephen Morris, Xiaosheng Mu, Mallesh Pai, Harry Di Pei, Jacopo Perego, Larry Samuelson, Vasiliki Skreta, Joel Sobel, Dale Stahl, Colin Stewart, Maxwell Stinchcombe, Rodrigo Velez, Thomas Wiseman, Muhamet Yildiz, and other seminar participants at UT, the 29th Stony Brook International Conference on Game Theory, 2018 Midwest Economic Theory Conference, 2018 Texas Economic Theory Camp, and 2018 SEA Conference for their excellent suggestions and helpful discussions.

1 Introduction

Individuals in positions of power, be they political leaders or the managers of firms, often suppress or censor bad news in order to improve their standing and prevent any threats to their authority. Such censorship is widely regarded to be undesirable. Nonetheless, we can imagine situations where the suppression of bad news may lead to better outcomes. For example, a political leader may be embarking on a radical reform that has the potential to be transformative. Being radical, the reform is also subject to teething troubles, and if the public were to become aware of all the difficulties, it might prematurely lose faith in the leader and replace her.

The decades of the 1960's - 1980's witnessed rapid industrialization and exceptionally high growth rates in many Asian developing countries and regions, including Singapore, South Korea, and Taiwan. This has been called an economic "miracle". Controversially, all these countries and regions were under authoritarian rule at the time. One hypothesis from a review by Sirowy and Inkeles (1990) links economic growth with authoritarianism: "the superior ability of an authoritarian regime to govern that facilitates economic growth is expressed indirectly by the social and political stability it fosters" (p. 130). This "political stability" may be sustained by different means, one of which is censorship.[1] As Rodan (2004) noticed, "[a]lmost by definition, authoritarian regimes involve censorship" (p. 1). On the other hand, we have many examples of leaders who persist with foolhardy projects, hiding all negative evidence. For instance, during China's Great Leap Forward (1958-1962), the central government's failure to access up-to-date local information due to local officials' concealment was partially responsible for the ensuing famine.[2] These considerations suggest that it is important to formally examine the implications of censorship in a dynamic context.

More specifically, this paper studies the interplay between learning and censorship in a dynamic environment. We consider a relationship between an agent who seeks to remain in office, and an evaluator who learns the competence of the agent from public news about the quality of the agent's project. The evaluator wants to retain the competent agent and dismiss the incompetent one, thereby terminating her project. The project of the agent gives rise to two news processes – good news and bad news. Good news is publicly observed, and confirms that the project is a good one. However, bad news can also arise, and is more likely when the project is a bad one. The agent can always suppress bad news when it materializes, but this is costly. Thus, the information that the evaluator receives is endogenously determined by the agent's censorship policy. We assume that the agent knows her own competence level, and consequently, competent agents and incompetent ones may well censor differently.

There are several applications of our model. The agent can be the manager of a

[1] Guriev and Treisman (2018) study how repression, co-option, propaganda, and censorship are used for an authoritarian regime to survive.

[2] See Li and Yang (2005).

division of a large firm, while the evaluator is the firm's overall manager or CEO. Alternatively, the agent might be an entrepreneur, with the evaluator being a venture capitalist who is funding the project. Finally, the agent might be a political leader, with the evaluator standing for the population. In these contexts, bad news can take several forms – mechanical breakdowns, reports of malpractice or customer complaints. Bad news can be suppressed – accounts can be "cooked", log files can be faked, and unhappy consumers could be mollified with refunds or gifts. None of these measures are costless; they take time and money, and psychological costs may be associated with dishonest behavior.[3] Similarly, politicians can arrest reporters, bribe witnesses, or shut down Internet forums, but this is also costly.

Basic economic intuition suggests that concealing information necessarily hurts the evaluator. Moreover, the possibility of censorship makes the evaluator more suspicious about the agent's performance. He does not know whether the reason that no bad news arrives is because the agent is competent, or because the agent is censoring. Thus the possibility of censorship also hurts a competent agent who has no way to prove that she was not hiding anything. It can only benefit an incompetent agent, since censorship helps her survive bad news. If the above intuition is correct, then the policy implication would be to reduce censorship by making it as hard and costly as possible.

However, this paper shows that the above intuition is only partially correct, and it crucially depends on the details of the information structure. Specifically, it depends on whether negative evidence is *conclusive*, i.e. it can only arise when the agent is incompetent. Indeed, when bad news can also arise when the agent is competent, we show that censorship can potentially increase the welfare of all parties, including the evaluator.

We now turn to the details of our model and its basic insights. All news is modeled as exponential/Poisson news. Good news can only arise for a competent agent, and is therefore conclusive. We consider two qualitatively different information structures. We first assume that bad news is also conclusive, i.e. it can only arise for an incompetent agent. We then consider the more general case of inconclusive bad news. The evolution of the evaluator's posterior belief about the competence of the agent in the absence of news depends on the agent's censorship policy, and on the (exogenous) parameters of the news processes. In our analysis, it will be useful to distinguish information structures according to whether the absence of news is "good" or "bad" news, i.e. whether the evaluator's posterior belief drifts up or down. We say that *good news arrives faster* when good news arrives faster than the incompetent agent's bad news. If good news arrives faster, then, regardless of the censoring decisions of the agent, the evaluator's belief is updated downwards in the absence of news. We say that *good news arrives slower*, when it arrives slower than the incompetent agent's bad news. In this case, the direction in which the evaluator's belief moves in the absence of news depends upon the censorship policy of the agent.

[3]See Rosenbaum, Billinger, and Stieglitz (2014) for a review on honesty experiments.

The simplest case is when bad news is conclusive, i.e. it only arises for the incompetent agent, and when good news arrives faster. Thus the public belief about the agent being competent will drift down in the absence of news regardless of the censorship policy, but censorship will accelerate the downward process. Our first observation is that the belief threshold at which the evaluator fires the agent is *independent* of the (incompetent) agent's censorship strategy, due to the fact that the option value of continuation depends only upon the possibility that good news arises. This observation allows us to use the logic of backwards induction to pin down behavior in any equilibrium. Since censorship is costly, the agent will incur the cost if and only if this secures her a sufficient increase in tenure. Thus, there is a unique belief threshold, at which the agent switches from *Full-Censorship* to *No-Censorship*. Obviously, the evaluator and the competent agent are worse off with censorship, which exactly confirms our initial intuition. However, the incompetent agent may also be worse off with censorship, since she has a shorter tenure, conditional on no bad news occurring, than when censorship is not possible. This happens whenever the Full-Censorship period is short. The benefit of censoring for the incompetent agent is that she survives bad news in the Full-Censorship period. When it is short, it means that the benefit of censoring is small, thus it will be overcome by the cost of censoring, i.e. the acceleration of the downward drifting of public belief.

Consider now the case where good news arrives faster, but where bad news is inconclusive, i.e. it can arise for both types of agent. Our main insight arises from the fact that the competent agent has greater incentives to censor bad news than the incompetent one. This is due to the fact that the competent agent knows that *good news may arise and secure her permanency in tenure*, which is not possible for the incompetent agent. Furthermore, the competent agent is also less likely to get further bad news. Therefore, the competent agent has a higher continuation value on the job, and the belief threshold at which she stops censoring, p^G, is lower than the threshold at which the incompetent agent stops censoring, $p^{B\dagger}$. Consequently, in the interval $[p^G, p^{B\dagger})$, bad news becomes conclusive *endogenously*; bad news can only come from the incompetent agent since the competent one always censors. By improving the quality of information, censorship can increase the payoff of the evaluator.[4]

Finally, let us consider the case where good news arrives slower. When bad news is conclusive, we cannot have a pure strategy equilibrium where the incompetent agent stops censoring at some threshold – if this were the case, the evaluator's belief would drift upwards, which would then make censoring bad news attractive. There is a critical threshold such that when this belief threshold is reached, both parties randomize. The incompetent agent randomizes between censoring and not, while the evaluator randomizes between firing the agent and retaining her.[5] Above this belief threshold,

[4]Indeed, for some parameter values it can also improve the payoffs of both types of agent, and therefore benefit all players.

[5]More precisely, in our continuous time model, the evaluator randomizes over stopping times.

the incompetent agent censors for sure. It is worth noting that the evaluator dismisses the agent at a higher threshold belief than he does in the absence of censorship, since the existence of censorship reduces the availability of future bad news, and since the option value of the evaluator also depends upon the arrival rate of bad news. This only happens when good news arrives slower, because the dismissal threshold depends critically on the arrival of bad news only if the belief updating process in the absence of news is upward, but not downward. We call the increase in the dismissal threshold belief the *discouragement effect* since censorship discourages the agent's incentive for learning. When bad news is inconclusive, we show that the very similar equilibrium still exists when the censoring cost is low. The evaluator and the incompetent agent use the same equilibrium strategy as in the conclusive bad news case, but the competent agent censors with probability 1 even when the belief is at the dismissal threshold. Basically, in this equilibrium, all bad news from both types of agent will be censored, except some bad news at the dismissal threshold. This hurts the evaluator, since her information quality is worsened. However, when the censoring cost is intermediate, there exists an equilibrium in which only the competent agent finds it optimal to censor when the belief is sufficiently high, but the incompetent agent never censors since the cost is too high for her. Again, this separation improves the quality of information and benefits the evaluator. We also find an *encouragement effect*; that is, the evaluator's dismissal threshold belief decreases since his incentive for learning is encouraged.

Our model has implications for how we should interpret censorship, and institutional measures against it. Let us interpret the cost of censorship as reflecting institutional structures against it. In some circumstances, i.e. when bad news only arises for the incompetent agent, we find that censorship is unambiguously bad, and thus the cost of censorship should be as high as possible. However, when bad news also arises for the competent agent, the evaluator prefers neither a very strong institution that prevents censorship, nor a very weak institution that allows for too much censorship. Mild censorship may be better than the two extreme cases.

The rest of the paper is organized as follows. Section 2 discusses the related literature. Section 3 introduces the model. Section 4 and Section 5 characterize equilibria and discuss the welfare effect for the conclusive bad news case and the inconclusive bad news case, respectively. Section 6 concludes. The Appendix provides all proofs.

2 Related Literature

This paper studies censorship in a dynamic environment of learning, where multiple signals arrives gradually. It mainly relates to four strands of literature.

First, it closely relates to a small literature on political censorship. Shadmehr and Bernhardt (2015) study a game between a ruler and a representative citizen, in which the ruler can censor a bad media report at a cost in order to mitigate the likelihood of revolution. Their paper, as well as this paper, assumes the role of the media is passive. Thus the focus is solely on the relation between the ruler and the citizen. They find that at the ex ante stage, before the ruler knows her type, she can increase her expected payoff by committing to censoring slightly less than she does in equilibrium where such a commitment power does not exist. One of our results also demonstrates that not only the agent before she knows her type, but also the agent who knows her type is bad can benefit from committing to no censorship.

Besley and Prat (2006) study the media capture problem. In their model, a media outlet maximizes his profits either from his audience who are interested in the informative news, or from a bad government who bribes him to keep silent about the bad news. Their paper focuses on the role of media outlets, and explicitly models the censoring cost as a direct or indirect transfer to a media outlet that is willing to forgo its readership and take a bribe. Gehlbach and Sonin (2014) and Eraslan and Ozerturk (2017) also examine the role of media outlets. The former takes a Bayesian persuasion approach as in Kamenica and Gentzkow (2011) such that a government can commit to an editorial policy for news release, while the latter studies a media outlet's reputation concerns that arise due to an information gatekeeping policy.

Egorov, Guriev, and Sonin (2009) consider the role of media outlets in monitoring bureaucrats.[6] In their model, a dictator needs a bureaucrat to implement her policy, but faces a moral hazard problem. To incentivize the bureaucrat to work, the dictator could rely only on the media report. However, a free media that solves the monitoring problem also exposes the dictator's incompetence to the public. Thus, if the dictator's income mainly depends on the bureaucrat's performance, she would rather risk being overthrown but allow a free media to monitor the bureaucrat. Lorentzen (2014) studies a regime change game in which a regime trades off the risk being overthrown and an informative media capable of monitoring the lower-level officials.

Edmond (2013) and Redlicki (2017) study other types of information manipulation in politics. Specifically, they study a global game played by citizens to attack a regime, in which the private signal of the citizens can be manipulated at a cost by the regime, either through shifting the mean of the signal in Edmond (2013), or through increasing the noise of the signal in Redlicki (2017). Guriev and Treisman (2018) study other means – propaganda, censorship, co-optation, and repression – that a dictator can employ to survive.

Although almost all papers on censorship consider one-shot interaction in a political setting, a notable exception is Smirnov and Starkov (2018).[7] They study censorship in

[6] Egorov, Guriev, and Sonin (2009) also extend their one-shot censorship model into a dynamic setting by assuming the dictator faces a stationary environment where censoring public bad news can effectively keep the ruler in power, and obtain a similar result as in their static model. In contrast, our model analyzes the (non-stationary) dynamics of the public confidence about the competence of the agent.

[7] Hauser (2017) also studies costly information manipulation in a dynamic environment. Different

product reviews. The main difference with our setup is that they assume censorship is costless, and some of the myopic consumers are naive, i.e. they do not understand that reviews can be censored. In contrast, we examine censorship with a strategic and forward-looking evaluator, and study the interdependence of censorship and learning, where costs play a crucial role.

Second, another related literature is on disclosure of verifiable information, beginning with Grossman (1981) and Milgrom (1981). A subset of the literature following Dye (1985) and Jung and Kwon (1988) studies the scenario where a receiver is uncertain about whether a sender possesses a piece of evidence. This is also the case in our setup since the evaluator's lack of news could happen either due to censorship or because no news has arrived. While the early literature exclusively focuses on static models, later papers extend them into different dynamic settings where multiple signals may arise in multiple periods, e.g. Shin (2003), Grubb (2011), Acharya, DeMarzo, and Kremer (2011), Guttman, Kremer, and Skrzypacz (2014).

Although the literature on disclosure games is large, very little attention has been paid to the case in which concealing information is costly. Three recent papers explore different implications of such a cost, all of which are very different from those that we focus on. Dye (2017) studies a static model of voluntary disclosure, in which a seller who has withheld information may be caught by a fact finder after the sale of an asset. In such an event, she has to make a damages payment to the buyer, the amount of which equals the product of the buyer's overpayment and a "damages multiplier". An important implication of this particular cost structure is that the seller would withhold more information as the punishment of the withholding goes up. Daughety and Reinganum (2018) study the problem of suppression of exculpatory evidence in prosecutions. In their model, a prosecutor wants to convict a defendant but also incurs a moral cost if she convicts an innocent defendant. The prosecutor also receives a penalty if she is caught for suppressing evidence by a reviewing judge. They extend their model to incorporate the teamwork of two prosecutors and show that this results in the concentration of authority regarding suppressing evidence. Kartik, Lee, and Suen (2017) provide a result on Bayesian updating, and apply it in a multi-sender disclosure game. They show that competition leads to more disclosure in the presence of a cost of concealment.

Third, a recent literature on dynamic information design is also related, e.g. Ely (2017), Renault, Solan, and Vieille (2017), Che and Hörner (2018). Their papers, as well as ours, study information manipulation in a dynamic environment. Their models rely on the principal's commitment power to design a flexible information disclosure policy to induce an agent to choose a desirable action for the principal, and also assume manipulating information is costless. In contrast, we focus on one particular kind of information manipulation – censorship. In addition, we do not assume the commitment

from the censorship literature that focuses on the ex post information suppression, he explores the case where ex ante effort can be exerted to slow down the arrival of bad news.

power, and assume censorship is costly.

Finally, this paper relates to two-armed Poisson bandit models in continuous time, e.g. Presman (1991), Keller, Rady, and Cripps (2005), Keller and Rady (2010, 2015). The evaluator here faces a two-armed bandit problem, in which the information generated by the risky arm is endogenous and partially controlled by the agent. We will borrow results from this literature to solve for the benchmark case in the absence of censorship.

3 The model

Two risk neutral players, an agent (she) and an evaluator (he), play a game in continuous time $t \in [0, \infty)$. Their discount rates are $\rho_0 > 0$ and $\rho_1 > 0$, respectively.

At time $t = 0$, nature chooses the type of the agent θ and the type of the project α, both from $\Theta = \{G, B\}$. We assume that $\theta = \alpha$, so that nature's choices are perfectly correlated. We assume that the agent observes nature's choice, while the evaluator does not. Let $p_0 \in (0, 1)$ denote the probability that nature chooses a type G project/agent.

The agent enjoys a flow payoff $w > 0$ that is independent with her type while she stays in her job, and has no payoffs after she is dismissed by the evaluator. Only the type G agent can succeed in her project. A success is publicly observable, and it arrives according to a Poisson process $\boldsymbol{S} = \{S_t\}_{t \geq 0}$ with an arrival rate $\gamma > 0$. The first success reveals that the type of the agent is G.

A success yields a lump-sum payoff $k > 0$ to the evaluator. The evaluator can choose a time $t \geq 0$ to irreversibly dismiss the agent. After the dismissal, the game ends, and the evaluator receives his outside option, the present discounted value of which is normalized to m.[8] We assume $h := \gamma k > m > 0$. Thus, the evaluator prefers the type G agent to the outside option, and prefers the outside option to the type B agent.

The project also potentially produces adverse news events, such as breakdowns. The arrival rate of such news depends upon the type of the project. If the project's type is G, it generates a piece of news at each jumping time of a Poisson process $\boldsymbol{N}^G = \{N_t^G\}_{t \geq 0}$ with an arrival rate $\beta^G \geq 0$. Conditional on the project's type being G, the success process $\boldsymbol{S}$ and the news process $\boldsymbol{N}^G$ are independent. If the project's type is B, it generates a piece of news at each jumping time of a Poisson process $\boldsymbol{N}^B = \{N_t^B\}_{t \geq 0}$ with an arrival rate β^B, where $\beta^B > \beta^G$. We call a piece of such news *bad news* since it happens more often to the type B agent, though it has no payoff consequence. We also call a success *good news*.

[8] It can also be interpreted as the flow payoff from the outside option, since they are equal after normalization (i.e. $\int_0^\infty \rho_1 e^{-\rho_1 t} m \, \mathrm{d}t = m$).

The agent can observe the bad news process. When a piece of bad news arrives, she can incur a lump-sum cost $c > 0$ to censor it. The evaluator can observe bad news if and only if the agent does not censor it, but he cannot distinguish whether a piece of bad news comes from $\boldsymbol{N}^G$ or $\boldsymbol{N}^B$, unless $\beta^G = 0$. Let $X_t^\theta \in \{1, 0\}$ be the type $\theta \in \Theta$ agent's censoring decision at time $t \geq 0$ when a piece of bad news has arrived at time t; $X_t^\theta = 1$ denotes censoring. A piece of bad news is *revealed* to the evaluator if it is not censored by the agent.

Histories and Strategies – Some care must be taken while defining the agent's and the evaluator's decision nodes in the game. Before the game ends,[9] the agent needs to make her censoring decisions only at dates when a piece of bad news arrives. On the other hand, the evaluator can dismiss the agent at any time.

At time t, a private history of the type $\theta \in \Theta$ agent is denoted by h_t^θ. It consists of a finite sequence of news realizations before and including date t, and her censoring decision for those bad news realizations before but not including date t. Let $\bar{h}_t^\theta$ be a typical private history for her at t when a piece of bad news has just arrived at t. Her strategy specifies a censoring probability $x_t^\theta \in [0, 1]$ at time t for each history $\bar{h}_t^\theta$; this strategy $\boldsymbol{x}^\theta = \{x_t^\theta\}_{t \geq 0}$ is progressively measurable with respect to the filtration induced by those histories. Additional requirements on the strategies will be imposed later to ensure that public beliefs are well-defined.

The evaluator only observes the censored news processes (i.e. the public history). At time t, a public history h_t is a finite sequence of news realizations that have not been censored before and including date t. Let $\mathscr{F} = \{\mathscr{F}_t\}_{t \geq 0}$ be the filtration induced by those public histories. The evaluator's pure strategy at time t for a public history h_t is a $\mathscr{F}$-stopping time T^t,[10] at which time he dismisses the agent.

Pure strategies are not sufficient to study equilibria. Thus, we introduce mixed strategies, which are cumulative distribution functions over stopping times. To be specific, a mixed strategy for the evaluator at time t for a public history h_t is a $\mathscr{F}$-adapted process $\boldsymbol{r}^t = \{r_\nu^t\}_{\nu \geq t}$, such that pathwise,

a) r_ν^t is non-decreasing and right continuous in $\nu \geq t$, and takes values in $[0, 1]$;

b) for any time $t' \geq t$, and $\nu \geq t'$, $r_\nu^{t'}$ is related to r_ν^t as follows:

$$r_\nu^t = r_{t'-}^t + (1 - r_{t'-}^t) r_\nu^{t'}, \quad \text{if} \quad r_{t'-}^t < 1,$$

where $r_{t'-}^t = \lim_{s \uparrow t'} r_s^t$ for $s > t$, and $r_{t-}^t = 0$.

a) requires r^t to be a cumulative distribution function over the stopping times T^t. b) requires r^t to be time consistent,[11] but imposes no restriction for $r_\nu^{t'}$ if $r_{t'-}^t = 1$. We

[9] All strategies are defined conditional on the game has not yet ended.

[10] The superscript t means the strategy is defined for a history at time t. This also applies to the strategy $\boldsymbol{r}^t$ which will be defined later.

[11] See Laraki, Solan, and Vieille (2005) and Riedel and Steg (2017) for details.

also call $\frac{\mathrm{d}r^t_\nu/\,\mathrm{d}\nu}{1-r^t_\nu}$ the (instantaneous) hazard rate whenever it exists.

Beliefs – The public belief at time t about the agent's type being G is $p_t = \mathbb{P}[\theta = G|\mathscr{F}_t]$, where the probability measure is induced by the public conjectured strategy $\tilde{\boldsymbol{x}}^\theta = \{\tilde{x}^\theta_t\}_{t\geq 0}$ of the type θ agent. In equilibrium, we must have $\tilde{x}^\theta_t = \mathbb{E}[x^\theta_t|h_t]$ for any t and and h_t. When a success arrives, the public belief jumps to 1. If at time t a piece of bad news is revealed to the evaluator, the public belief jumps from $p_{t\text{-}}$ to

$$J(p_{t\text{-}}, \tilde{x}^G_t, \tilde{x}^B_t) := \frac{p_{t\text{-}}\beta^G(1-\tilde{x}^G_t)}{p_{t\text{-}}\beta^G(1-\tilde{x}^G_t) + (1-p_{t\text{-}})\beta^B(1-\tilde{x}^B_t)},$$

when the denominator is positive. Observe that in the absence of censorship (i.e. $\tilde{x}^G_t = \tilde{x}^B_t = 0$), a piece of bad news makes the public belief jump from $p_{t\text{-}}$ down to $j(p_{t\text{-}}) := J(p_{t\text{-}}, 0, 0) < p_{t\text{-}}$.

If no news is revealed to the evaluator over the time interval $[t, t+\mathrm{d}t)$, the public belief p_t evolves continuously according to the following ordinary differential equation (ODE),

$$\dot{p}_t = g(p_t, \tilde{x}^G_t, \tilde{x}^B_t) := -p_t(1-p_t)[\gamma + (1-\tilde{x}^G_t)\beta^G - (1-\tilde{x}^B_t)\beta^B]. \tag{ODE}$$

To make sure the above (ODE) with initial condition $p_0 \in (0,1)$ admits a well-defined solution, we require $\tilde{\boldsymbol{x}}^\theta$ to satisfy the following requirement. For $\theta \in \Theta$, we say $\tilde{\boldsymbol{x}}^\theta$ is *admissible* if for any public history $h_{t'}$ and $t' > 0$, $\{\tilde{x}^\theta_t\}_{t\leq t'}$ is piecewise continuous in t with finite cutoffs, and it is right continuous with left limits at the cutoffs. Given any admissible $\tilde{\boldsymbol{x}}^G$ and $\tilde{\boldsymbol{x}}^B$, the ODE with initial condition $p_0 \in (0,1)$ admits a unique solution (or trajectory) $\boldsymbol{p} = \{p_t\}_{t\geq 0}$ which is piecewise continuously differentiable.[12]

Payoffs – For $\bar{h}^\theta_t$, given $\boldsymbol{r}^t$ and $\boldsymbol{x}^\theta_t = \{x^\theta_\nu\}_{\nu\geq t}$, the type θ agent's expected payoff at time t is

$$v^{\boldsymbol{r},\boldsymbol{x}^\theta}_\theta(\bar{h}^\theta_t) = \mathbb{E}\left[-\rho_0\; c\; X^\theta_t + \int_t^{T^t} \rho_0\; e^{-\rho_0(\nu-t)}(w\,\mathrm{d}\nu - c\; X^\theta_\nu\,\mathrm{d}N^\theta_\nu)\middle|\; \bar{h}^\theta_t\right],$$

where the expectation is taking over T^t, $\{X^\theta_\nu\}_{\nu\geq t}$, and $\{N^\theta_\nu\}_{\nu\geq t}$.

For h_t, given $\boldsymbol{r}^t$ and $\tilde{\boldsymbol{x}}^\theta_t = \{\tilde{x}^\theta_\nu\}_{\nu\geq t}$, the evaluator's expected payoff at time t is

$$u^{\boldsymbol{r},\tilde{\boldsymbol{x}}}(h_t) = \mathbb{E}\left[\int_t^{T^t} \rho_1\; e^{-\rho_1(\nu-t)}\;\mathbb{1}_{\theta=G}\,k\,\mathrm{d}S_\nu + e^{-\rho_1(T^t-t)}\; m\middle|\; h_t\right],$$

where the superscript $\tilde{\boldsymbol{x}} := \{\tilde{\boldsymbol{x}}^G, \tilde{\boldsymbol{x}}^B\}$, and the expectation is taking over T^t, $\{S_\nu\}_{\nu\geq t}$, $\{\tilde{X}^\theta_\nu\}_{\nu\geq t}$, and $\{N^\theta_\nu\}_{\nu\geq t}$.

[12]Given piecewise continuous $\tilde{\boldsymbol{x}}^\theta$, it is easy to see that $g(p, \tilde{x}^G_t, \tilde{x}^B_t)$ and $g_p(p, \tilde{x}^G_t, \tilde{x}^B_t)$ are piecewise continuous in t for a fixed p, and $g(p, \tilde{x}^G_t, \tilde{x}^B_t)$ is continuously differentiable in p for a fixed t. Thus, the standard local existence and uniqueness results for ODE ensure a unique solution in the interval where both $\tilde{x}^G_t$ and $\tilde{x}^B_t$ are continuous. Also, a unique global solution can be obtained through concatenation.

Let $V_\theta^{\boldsymbol{r}}(\bar{h}_t^\theta) = \sup_{\{x_\nu^\theta\}_{\nu\geq t}} v_\theta^{\boldsymbol{r},\boldsymbol{x}^\theta}(\bar{h}_t^\theta)$ be the value function of the type θ agent, and $U^{\tilde{\boldsymbol{x}}}(h_t) = \sup_{\{r_\nu^t\}_{\nu\geq t}} u^{\boldsymbol{r},\tilde{\boldsymbol{x}}}(h_t)$ be the value function of the evaluator, given the strategies of other players.

Equilibria – We study Perfect Bayesian Equilibria (PBE), which consist of strategies $\boldsymbol{x}^\theta = \{x_t^\theta\}_{t\geq 0}$ and $\boldsymbol{r} = \{r_\nu^t\}_{t\geq 0,\nu\geq t}$, a public belief $\boldsymbol{p} = \{p_t\}_{t\geq 0}$, and the conjectured strategy $\tilde{\boldsymbol{x}}^\theta = \{\tilde{x}_t^\theta\}_{t\geq 0}$ such that

1. For each $\bar{h}_t^\theta$, $\boldsymbol{x}_t^\theta$ is optimal for the type θ agent, given $\boldsymbol{p}$ and $\boldsymbol{r}$;
2. For each h_t, $\boldsymbol{r}^t$ is is optimal for the evaluator, given $\boldsymbol{p}$ and $\tilde{\boldsymbol{x}}^\theta$;
3. $\boldsymbol{p}$ is updated according to $\tilde{\boldsymbol{x}}^\theta$ and Bayes rule whenever possible, and $\tilde{\boldsymbol{x}}^\theta = \mathbb{E}[\boldsymbol{x}^\theta|\mathscr{F}]$.

It is well known in the two-armed bandit literature that the optimal allocation rule is a cutoff rule with respect to the public belief. In our setup, cutoff strategies are of particular interest. We say the evaluator's strategy is a pure cutoff strategy with a cutoff public belief $\hat{p} \in (0,1)$ if it is a pure strategy $T^t = \inf\{\nu \geq t : p_\nu \leq \hat{p}\}$ for any public history h_t. This means that the evaluator dismisses the agent whenever the public belief is below or equal to the cutoff $\hat{p}$. The evaluator's strategy is a mixed cutoff strategy with a cutoff $\hat{p} \in (0,1)$ if for any public history h_t and $\nu \geq t$,

$$r_\nu^t = \begin{cases} 1 & \exists s \in [t,\nu], p_s < \hat{p}, \\ 0 & \forall s \in [t,\nu], p_s > \hat{p}. \end{cases}$$

This means that the evaluator never dismisses the agent when the public belief is above the cutoff $\hat{p}$, but dismisses the agent immediately when the belief is below $\hat{p}$.

Similarly, we say the type θ agent's strategy is a pure (resp. mixed) cutoff strategy with a cutoff $\hat{p} \in (0,1)$ if for any history $\bar{h}_t^\theta$,

$$x_t^\theta = \begin{cases} 1 & \text{if } p_t \in (\hat{p}, 1), \\ 0 & \text{if } p_t \leq \hat{p} \text{ (resp. if } p_t < \hat{p}). \end{cases}$$

This means that the agent censors bad news when the public belief is above the cutoff $\hat{p}$ (but below 1), and she stops censoring when the belief is below $\hat{p}$. The difference between the pure and the mixed strategies is that the mixed strategy allows mixing at the cutoff belief.[13] We call a PBE a *cutoff equilibrium* if the evaluator uses a cutoff strategy.

[13] A mixed cutoff strategy first needs to satisfy the definition of a strategy; thus, the mixing is also a contingent plan.

4 Conclusive bad news

We first study the baseline model when the type G agent does not have any bad news (i.e. $\beta^G = 0$). Thus, the type G agent is a passive player who does not have any action. A piece of revealed bad news is conclusive evidence to the evaluator that the agent's type is B. In this case, we will show that censorship not only always makes the evaluator and the type G agent worse off, but also sometimes makes the type B agent worse off.

Now a revealed signal (good or bad news) resolves all uncertainty about the type of the agent. Hence, after seeing a piece of good news, the evaluator never dismisses the agent. After seeing a piece of revealed bad news, the evaluator dismisses the agent immediately. It follows that the type B agent will not censor any bad news after a piece of bad news has already been revealed to the evaluator, even if the evaluator has deviated and has not yet dismissed her. Therefore, what remains to be determined is simply (1) the evaluator's strategy under the public history at time t (denoted by $\varnothing_t$) where no signal has been revealed to the evaluator, and (2) the type B agent's strategy under her private history at time t (denoted by $\bar{\varnothing}_t^B$) where no signal has been revealed to the evaluator as yet and a piece of bad news has just arrived. We say the evaluator experiments when he does not dismiss the agent under the public history $\varnothing_t$.

4.1 No censorship when the censoring cost is high

A prohibitively high cost will prevent censorship. This section determines the threshold cost above which censorship does not exist in equilibrium.

Observe that it is strictly dominant for the evaluator to dismiss the agent when he sees bad news. Given this restriction, the best possible strategy of the evaluator for the type B agent is one where she is only dismissed when bad news arises. We use $\boldsymbol{r}_\infty$ to denote this strategy. Thus, the dismissal time is the first time that the type B agent does not censor a piece of bad news when it arrives. Given this strategy, the type B agent faces a stationary problem: consider the type B agent's problem at the history $\bar{\varnothing}_t^B$. The agent can choose between censoring or not. Her value function satisfies the following Bellman equation,

$$V_B^{\boldsymbol{r}_\infty}(\bar{\varnothing}_t^B) = \max\left\{\mathbb{E}[-\rho_0\, c + \int_0^{\tau^B} \rho_0\, e^{-\rho_0 \nu} w \,\mathrm{d}\nu + e^{-\rho_0 \tau^B} V_B^{\boldsymbol{r}_\infty}(\bar{\varnothing}_{\tau^B}^B)], 0\right\},$$

where τ^B is the arrival time of the next piece of bad news from the type B agent. She will be dismissed and obtain 0 payoff if she does not censor this piece of bad news. Otherwise, by censoring it, she incurs a cost $\rho_0 c$, but can stay in her job until the next piece of bad news arrives at τ^B, and enjoy her continuation value after τ^B. Clearly, $V_B^{\boldsymbol{r}_\infty}(\bar{\varnothing}_t^B)$ does not depend on t. Thus we use $V_B^{\boldsymbol{r}_\infty}(\bar{\varnothing}^B) := V_B^{\boldsymbol{r}_\infty}(\bar{\varnothing}_t^B)$.

It is easy to check that

$$V_B^{\boldsymbol{r}_\infty}(\bar{\varnothing}^B) = \max\left\{w - (\rho_0 + \beta^B)c, 0\right\}.$$

Hence, the type B agent has a strict incentive not to censor any bad news if and only if $c > \underline{c}$, where $\underline{c} = \frac{w}{\rho_0+\beta^B}$. Since any equilibrium strategy of the evaluator gives a weakly lower continuation payoff to the type B agent than $\boldsymbol{r}_\infty$, she will never censor any bad news when $c > \underline{c}$. Consequently, the evaluator faces a decision problem in the absence of censorship. His optimal policy is a cutoff strategy in which the cutoff belief is determined by the evaluator's trade-off between exploitation and exploration, as demonstrated in the bandit literature. The following proposition summarizes the result.

Proposition 1. *Assume $\beta^G = 0$ and $c > \underline{c}$. There is a unique PBE with the following features:*

- *Assume that $\gamma > \beta^B$. The type B agent never censors any bad news, and the evaluator dismisses the agent whenever the public belief is below $p_{fast} = \frac{\rho_1 m}{\rho_1 h + \gamma(h-m)}$.*
- *Assume that $\gamma < \beta^B$. The type B agent never censors any bad news, and the evaluator dismisses the agent whenever the public belief is below $p_{slow} = \frac{\rho_1 m}{\rho_1 h + \beta^B(h-m)}$.*

Notice that we have made a distinction about what kind of news arrives faster. We say *good news arrives faster* if it arrives faster than bad news from the type B agent (i.e. $\gamma > \beta^B$), and *good news arrives slower* if it arrives slower than bad news from the type B agent (i.e. $\gamma < \beta^B$). This distinction is often made in Poisson bandit models where news arrives according to exogenous Poisson processes. It determines the direction of the evolution of the public belief in the absence of news. In a Poisson bandit problem, if good news arrives faster, then the public belief drifts up in the absence of news, and the cutoff belief in the evaluator's optimal strategy is p_{fast}. This is similar to the breakthrough case as in Keller, Rady, and Cripps (2005). If good news arrives slower, then the public belief drifts down in the absence of news, and the cutoff belief in the evaluator's optimal strategy is p_{slow}. This is similar to the breakdown case as in Keller and Rady (2015). We will see later that this distinction is also useful for our analysis where the news process is endogenously determined by the equilibrium censorship strategy.

This equilibrium also demonstrates what would happen if censorship were not possible, i.e. the *no censorship benchmark* (NCB). To analyze the welfare effect of censorship later in the paper, we often compare a player's equilibrium payoff with his or her expected payoff in the NCB.

For the rest of the baseline model (conclusive bad news), we will assume that the censoring cost is not prohibitively high: $0 < c < \underline{c}$. We again distinguish two cases, according to the direction of drift of the posterior belief in the absence of news and in the absence of censorship.

4.2 Faster good news

4.2.1 Equilibrium

First we consider the case where good news arrives faster (i.e. $\gamma > \beta^B$). For such a project, (ODE) indicates that the public belief drifts down in the absence of news, regardless of the agent's censorship strategy. The greater the intensity of censorship, the faster the public belief declines. However, the evaluator's optimal strategy (as a function of the public belief) does not depend on how much information has been censored, as illustrated in Figure 1.

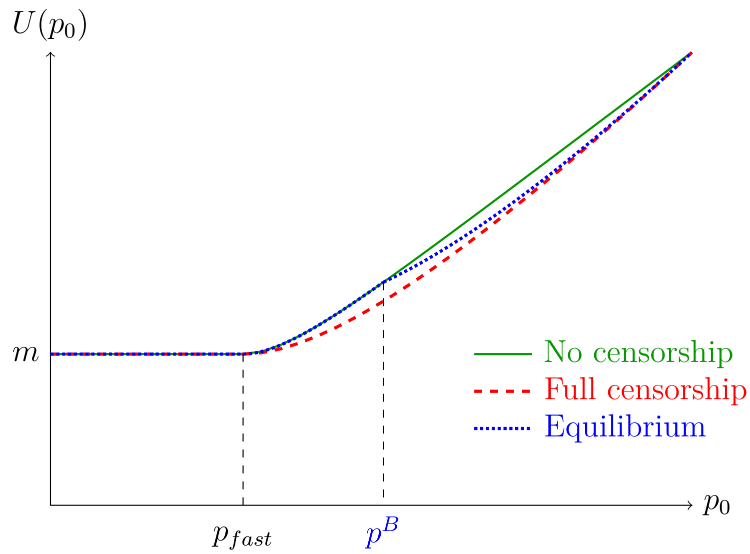

Figure 1: The evaluator's value functions (conclusive bad news & faster good news)

Figure 1 shows the evaluator's value functions under different censorship policies. The horizontal axis is the prior belief. Consider two extreme censorship policies – no censorship (green solid curve) where the agent never censors bad news, and full censorship (red dashed curve) where the agent always censors bad news. The no censorship policy gives the evaluator the highest payoff for any prior belief, while the full censorship policy gives him the lowest payoff. The evaluator's best response as a function of the public belief, however, is actually the same for those two extreme censorship policies, i.e. he dismisses the agent if and only if the public belief is below p_{fast}. This is because the evaluator learns from good news at the margin in both cases – he will dismiss the agent if no news is revealed or if a piece of bad news is revealed, so only good news can change his action. Hence, the optimal cutoff belief p_{fast} does not depend on the arrival of bad news, and therefore is the same for the two extreme cases. From the evaluator's perspective, the informativeness of any censorship strategy of the agent lies between those two extreme strategies, in term of the Blackwell order, at each instant.

Thus, for any admissible strategy of the agent, the evaluator's value must be in between his values when he faces those two extreme censorship policies. Hence, he will use the same cutoff strategy. The blue dotted curve in Figure 1 illustrates the evaluator's value function when the type B agent uses a pure cutoff strategy with the cutoff p^B. We will show later in Proposition 2 that this strategy is also the evaluator's equilibrium strategy. We should keep in mind that although the evaluator uses the same cutoff strategy, the greater the agent censors, the faster the same cutoff belief is reached in the absence of news.

Given the evaluator's strategy, we can pin down the type B agent's strategy and the unique PBE.

Proposition 2. *Assume $\beta^G = 0$, $c < \underline{c}$ and $\gamma > \beta^B$. There exists a unique PBE. The equilibrium features a pair of pure cutoff strategies.*

1. *The evaluator dismisses the agent whenever the public belief is below p_{fast};*
2. *The type B agent censors bad news whenever the public belief is above p^B, where $p^B > p_{fast}$.*

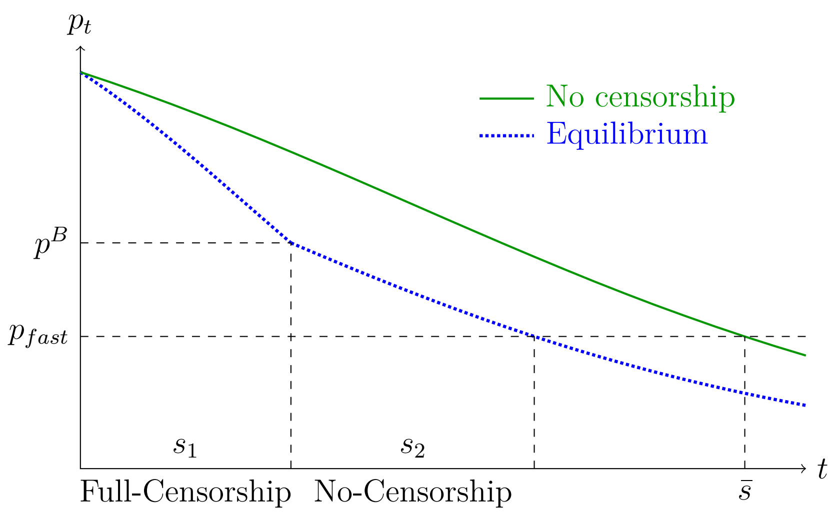

Figure 2: The belief evolution in the absence of news (conclusive bad news & faster good news)

Figure 2 shows the evolution of the public belief in the absence of news in the equilibrium (blue dotted curve). The evaluator will dismiss the agent whenever the public belief is below p_{fast}, which will happen eventually without the arrival of good news. Hence, the type B agent knows that she will be dismissed in the end. As the public belief approaches the threshold p_{fast}, she knows that she has little time left before she is fired. Her benefit of censoring – the extra time she is able to stay on the job – becomes relatively small, compared with the cost of censoring. Thus, the type B agent will give up censorship before the public belief reaches p_{fast}. In addition, before giving up censorship, the type B agent would censor all bad news with probability one. This is because the type B agent's continuation value after censoring is increasing in the public belief which is a decreasing function of time. Whenever it is weakly optimal to censor bad news at time t before the public belief drifts down to p_{fast}, it is strictly optimal to censor all bad news before time t. Hence, the censorship policy has two phases. At the beginning when the public belief is high (i.e. higher than p^B), the agent censors all bad news, we call it the Full-Censorship period; the length of this period is denoted by s_1 in Figure 2. After the Full-Censorship period when the public belief is low (i.e. lower than p^B), the agent stops censoring. We call this the No-Censorship period; the length of this period is stochastic since it is ended by the arrival of bad news, and the maximal length is denoted by s_2 in Figure 2. If the agent could commit to a policy of never censoring, then her maximal duration on the job would be $\bar{s}$ in Figure 2.

The following indifference condition determines the threshold at which the agent stops censoring:

$$\mathbb{E}[-\rho_0\ c + \int_0^{s_2 \wedge \tau^B} \rho_0\ e^{-\rho_0 \nu} w\, \mathrm{d}\nu] = 0,$$

where s_2 the maximal length of the No-Censorship period, and τ^B is the arrival time of the next piece of bad news from the type B agent. Thus, the expected continuation value in the No-Censorship period must be equal to the censoring cost. Clearly, the larger the censoring cost, the sooner it becomes not profitable to censor bad news, and the longer the No-Censorship period is. At last, s_2 determines the cutoff belief p^B, at which the agent switches from the Full-Censorship policy to the No-Censorship policy.

4.2.2 Welfare: Equilibrium versus NCB

We now consider the welfare consequence of censorship. Specifically, we compare two scenarios: players' equilibrium payoffs, and their expected payoffs in the absence of censorship (i.e. their payoffs in the NCB). Since the equilibrium strategy of the type B agent is to not censor any bad news when $p_t \leq p^B$, nothing would be different between these two scenarios. Thus, let us only consider $p_0 > p^B$.

The evaluator is worse off in equilibrium, compared with the NCB.[14] Figure 1 illustrates the evaluator's welfare loss due to censorship. This is because censorship garbles information. No news in equilibrium means either nothing really happened, or bad news has arrived but has been censored. Hence, the public belief drifts down faster when no news is revealed, compared with the NCB. The type G agent is also worse off

[14] We always compare the equilibrium payoff with the payoff in the NCB, unless stated otherwise. Hence, we will not repeat it every time.

in equilibrium, since, in order to stay in her job, she now has to have a success within a shorter period of time before the public belief drifts down below the evaluator's cutoff belief.

It is worth noticing that the type B agent may also be worse off in equilibrium. Since the type B agent can survive bad news when she censors in the Full-Censorship period, censorship does provide her a higher expected flow payoff. This is her benefit from censorship. However, the cost is that the possibility of censorship also drives the public belief down faster in the absence of news, compared with the NCB. Another way to look at this is the martingale property of the public belief. Since bad news would be censored, the fact that the public belief cannot jump to 0 means that the drifting down process must be accelerated. We can show that whenever the actual censoring period (i.e. the Full-Censorship period) is short, the benefit from censorship is dominated by the cost. Hence, we have the following result.

Proposition 3. *Assume $\beta^G = 0$, $c < \underline{c}$, $\gamma > \beta^B$ and $p_0 > p^B$. There exists a $s_1^* > 0$ such that the type B agent is worse off in equilibrium than she is in the NCB if and only if $s_1 < s_1^*$.*

The Full-Censorship period is an equilibrium object, but we can express this result in terms of primitive variables. The most straightforward comparative statics is for the prior belief. When the prior belief is high, the Full-Censorship period is long, thus the type B agent is better off in equilibrium. When the prior belief is low but higher than p^B, it means the Full-Censorship period is short, thus the type B agent is worse off in equilibrium.[15]

Another interesting comparative statics is for the censoring cost. The censoring cost stands for how hard it is to hide evidence from the evaluator. It also represents the strength of anti-censorship institutions. The cost directly determines when the agent is willing to give up censorship (i.e. the threshold belief p^B) and the length of the No-Censorship period. Thus, it also indirectly determines how long the agent censors bad news (i.e. the length of the Full-Censorship period). For a given prior belief, if the censoring cost is very high (but still below $\underline{c}$), it means the No-Censorship period is very long such that the Full-Censorship period disappears. Then, obviously, the equilibrium payoff is the same as the payoff in the NCB. If the censoring cost is very low, it means the No-Censorship period is very short, which in turn means the Full-Censorship is long. If it is long enough, then the type B agent is better off in equilibrium. However, when the censoring cost is intermediate, then the Full-Censorship exists but is short, and the type B agent is worse off in equilibrium. The result is summarized below.

Corollary 1. *Assume $\beta^G = 0$, $c < \underline{c}$, $\gamma > \beta^B$ and $p_0 > p_{fast}$. There exists a $0 \leq \mathfrak{c}_1 <$*

[15]Similarly, if the value of the outside option m is low, then the type B agent is better off in equilibrium. If m is high but not too high such that the Full-Censorship period still exists, then the type B agent is worse off in equilibrium.

$\mathfrak{c}_3 < \bar{c}$, the type B agent is better off (resp. worse off) in the equilibrium than she is in the NCB when $c \in (0, \mathfrak{c}_1)$ (resp. when $c \in (\mathfrak{c}_1, \mathfrak{c}_3)$). In addition, $\mathfrak{c}_1 > 0$ if and only if $\bar{s} > s^*$ for some $s^* > 0$.

Another way to understand the censoring cost is to examine the dependence of the equilibrium payoff with respect to the cost, which is shown in the following result.

Corollary 2. *Assume $\beta^G = 0$, $c < \underline{c}$, $\gamma > \beta^B$ and $p_0 > p_{fast}$. There exists a $\mathfrak{c}_2 \in [\mathfrak{c}_1, \mathfrak{c}_3)$, the equilibrium payoff of the type B agent is decreasing (resp. increasing) in c when $c \in [0, \mathfrak{c}_2]$ (resp. when $c \in [\mathfrak{c}_2, \mathfrak{c}_3]$). In addition, there exists a $s^{**} \in (0, s^*)$, such that $\mathfrak{c}_2 > \mathfrak{c}_1 \geq 0$ (resp. $\mathfrak{c}_2 = \mathfrak{c}_1 = 0$) when $\bar{s} > s^{**}$ (resp. when $\bar{s} \leq s^{**}$).*

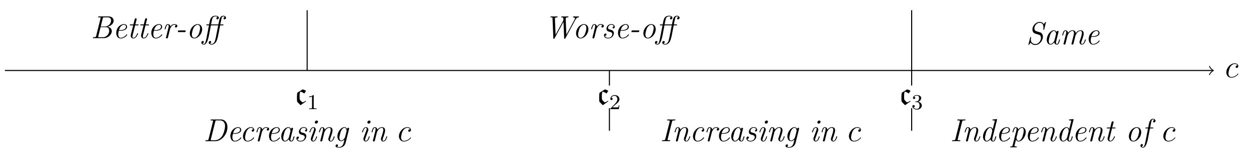

The censoring cost servers as a "commitment" device for the agent. An increase in the censoring cost decreases the type B agent's expected flow payoff in the Full-Censorship period since censoring is more costly, but it increases her payoff in the No-Censorship period. Moreover, an increase in the cost also increases the length of the No-Censorship period, but decreases the length of the Full-Censorship period.

As we illustrated in Corollary 1, when the censoring cost is very small (i.e. $c \in (0, \mathfrak{c}_1)$), the Full-Censorship period is relatively long, and the type B agent is better off in equilibrium. In this case, the negative effect of an increasing censoring cost would dominate the positive effect, since the expected flow payoff in the Full-Censorship period is more important. In this case, the type B agent's equilibrium payoff is decreasing in the censoring cost. This means that the type B agent who is better off with censorship must prefer a weaker institution (i.e. a smaller censoring cost).

The inverse statement is not true. When the censoring cost is high such that the Full-Censorship exists but is short (i.e. $c \in (\mathfrak{c}_1, \mathfrak{c}_3)$), the type B agent is worse off in equilibrium. However, it is not necessary for her to prefer a stronger institution. In general, her equilibrium payoff is not monotonic in the cost. Only when the cost is high enough (i.e. $c \in (\mathfrak{c}_2, \mathfrak{c}_3)$), the commitment (positive) effect of an increasing censoring cost dominates the negative effect. In this case, her equilibrium payoff is increasing in the censoring cost, and a stronger institution would be preferred by the type B agent.

4.3 Slower good news

4.3.1 Equilibrium

Now we turn to the other case where good news arrives lower (i.e. $\gamma < \beta^B$). In the NCB, since the evaluator would expect to see bad news more often, the public belief will drift up when no news arrives. Hence, "no news is good news", and the evaluator only dismisses the agent when a piece of bad news arrives, if he starts with a prior belief that is not too low ($p_0 > p_{slow}$).

However, the possibility of censorship changes everything. Foremost, the evaluator's optimal dismissal policy does depend on the agent's censorship policy, which is different from the case where good news arrives faster. In the NCB, the evaluator's optimal policy is to dismiss the agent whenever the public belief is below p_{slow}. However, if the agent censors all bad news, then the evaluator's optimal policy is to dismiss the agent when the public belief is below p_{fast}, as shown in the last section. Hence, censorship discourages the evaluator's incentive for exploring the competence of the agent. In addition, we also know from the above two extreme cases that the evaluator would never dismiss the agent when the public belief is higher than p_{fast}, and he would dismiss the agent immediately when the public belief is below p_{slow}, since again the evaluator's value function for any censorship policy is in between his values given by the two extreme censorship policies. This is illustrated in Figure 3.

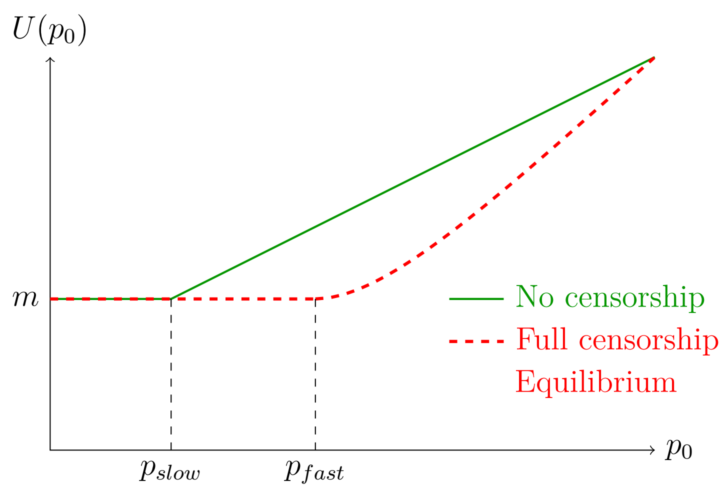

Figure 3: The evaluator's value functions (conclusive bad news & slower good news)

When the public belief is higher than p_{fast}, perhaps ironically, the fact that the type B agent can keep her job for sure also implies that she must censor all bad news in equilibrium when the public belief is above p_{fast}. First, when no news is revealed to the evaluator, the public belief must eventually drift down below p_{fast}, otherwise the type B agent can stay in her job forever by censoring all bad news, which in turn drives the public belief down – a contradiction. Second, the only way that the public belief drifting down below p_{fast} is that the agent censors bad news before the public is below p_{fast}. In addition, if the agent censors bad news at a public belief that is above p_{fast}, she must also find it optimal to censor bad news when the public belief is even higher, since a longer stay in her job gives her a higher incentive for censorship. Those together imply that the type B agent will censor all bad news when the public belief is above p_{fast}, and the only direction that the public belief drifts is downward, when it is above p_{fast}. Essentially, the evaluator faces the worst information, and would use a cutoff p_{fast} to dismiss the agent.

However, a pure strategy that dismisses the agent when the cutoff belief hits p_{fast} cannot be supported in equilibrium. If the evaluator uses such a pure strategy, he needs to dismiss the agent deterministically if no news arrives for some time, since the type B agent censors all bad news when the public belief is above p_{fast}. However, this means that the type B agent would like to give up censorship even when the public belief is above p_{fast}, if the evaluator would dismiss her deterministically in a short time. Such a contradiction implies that we need a mixed strategy for the evaluator. This also means that the type B agent must also mix at the cutoff belief p_{fast}, since censoring too much would induce the evaluator to use the pure strategy, while censoring too little would induce the public belief to drift up above p_{fast}. The former contradicts with the fact that the evaluator must uses a mixed strategy, and the latter contradicts with the fact that belief cannot drift up above p_{fast}.

Hence, the unique cutoff equilibrium is summarized in the following result.

Proposition 4. *Assume $\beta^G = 0$, $c < \underline{c}$ and $\gamma < \beta^B$. There exists a unique cutoff equilibrium, which features a pair of mixed cutoff strategies with a common cutoff p_{fast}.*

1. *The evaluator dismisses (resp. does not dismiss) the agent when the public belief is below (resp. above) p_{fast}. He uses a constant hazard rate $z^* := \frac{w}{c} - \rho_0 - \beta^B$ to dismiss the agent at the cutoff belief p_{fast}.*
2. *The type B agent censors (resp. does not censor) bad news when the public belief is above (resp. below) p_{fast}. She censors bad news with a constant probability $x^{B*} := 1 - \frac{\gamma}{\beta^B}$ at the cutoff belief p_{fast}.*

Figure 4 shows the evolution of the public belief in the absence of news in the equilibrium (red dashed curve). The censorship policy also has two phases. Since the type B agent censors all bad news when the public belief is above p_{fast}, we still have the Full-Censorship period at the beginning with length $\hat{s}$. When the public belief reaches p_{fast}, the type B agent censors bad news with a constant probability such that the arrival rate of good news equals to the arrival rate of the non-censored bad news (i.e. $\gamma = (1 - x^{B*})\beta^B$). Hence, the public belief would stay constant at p_{fast}, if no news arrives, which makes the problem becomes stationary. Moreover, the evaluator's constant hazard rate is also chosen to make the type B agent indifferent between censoring

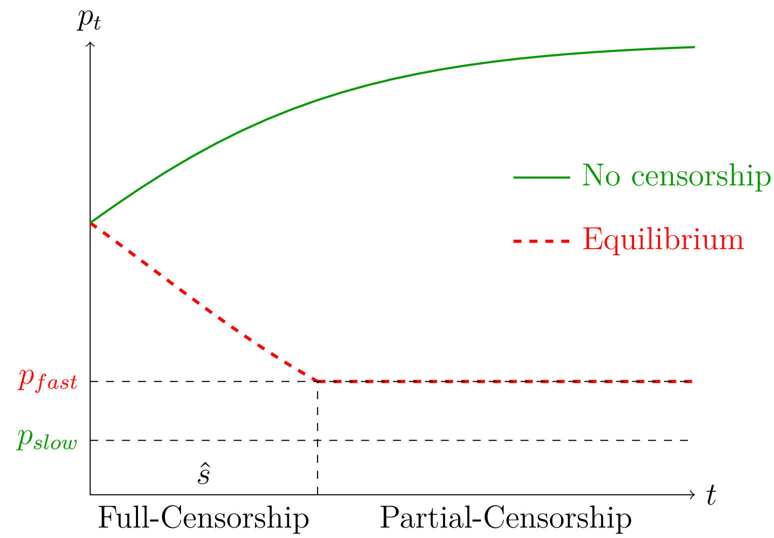

Figure 4: The belief evolution in the absence of news (conclusive bad news & slower good news)

and not. We call this period as the Partial-Censorship period.

This is the unique cutoff equilibrium. If the evaluator uses a cutoff strategy with a different cutoff belief $\hat{p}$, then it must be below p_{fast}, because the evaluator never dismisses the agent when the public belief is above p_{fast}. However, the exact argument which shows the type B agent censors all bad news when the public belief is above p_{fast} in equilibrium implies that now she must censor all bad news when the public belief is above $\hat{p}$. This means even less information is left for the evaluator. Hence, he does not want to use a cutoff strategy with a lower cutoff, since his incentive for learning is discouraged. Therefore, we obtain the unique cutoff equilibrium.

4.3.2 Welfare: Equilibrium versus NCB

We still compare players' payoffs in the cutoff equilibrium with their payoffs in the NCB. Clearly, the evaluator is worse off in equilibrium, as illustrated in Figure 3. Actually, this is the worst payoff he can obtain from any censorship policy. Equivalently, the evaluator cannot use any information from bad news; he can only rely on good news. That is why his incentive for learning is discouraged so that he increases the cutoff belief from p_{slow} to p_{fast} in his cutoff strategy. This is a new effect that does not appear when the good news arrives faster; we call it the *discouragement effect*. Due to this discouragement effect, when the prior belief is in between p_{slow} and p_{fast}, neither type of agent would have a chance to start her job in equilibrium, but she could at least have some positive time in her job in the NCB. Thus, when $p_0 \in (p_{slow}, p_{fast})$, both types of agent are worse off in equilibrium, compared with the NCB.

When the prior belief is above p_{fast}, the type G agent would stay in her job forever in the NCB, while she could be dismissed with a strictly positive probability in the equilibrium. Hence, she is worse off in equilibrium, compared with the NCB. For the type B agent, She has a higher payoff in the Full-Censorship period since she survives bad news, but a lower payoff in the Partial-Censorship period. Since her total payoff is the average payoff in those two periods, she is also worse off in equilibrium whenever the Full-Censorship period is short. We summarize the result below.

Proposition 5. *Assume $\beta^G = 0$, $c < \underline{c}$ and $\gamma < \beta^B$.*

1. *When $p_0 \in (p_{slow}, p_{fast}]$, the type B agent's is worse off in equilibrium than she is in the NCB;*
2. *When $p_0 > p_{fast}$, there exists a $\hat{s}^* > 0$ such that the type B agent's is worse off in equilibrium than she is in the NCB if and only if $\hat{s} < \hat{s}^*$.*

The Full-Censorship period, again, is an equilibrium object, but we can easily express the condition in terms of primitive variables. For example, when the prior belief is high, it means the Full-Censorship period is long, thus the type B agent is better off in equilibrium. When it is low but still higher than p_{fast}, it means the Full-Censorship period exists but is short, thus the type B agent is worse off in equilibrium.

For the censoring cost, different from the faster good news case, it does not determine the length of the Full-Censorship period, since the type B agent switches her strategy to partial censorship at a particular belief p_{fast} that does not depend on the cost. Hence, whether the type B agent is better off in equilibrium does not depend on the cost. However, an increase in the censoring cost decreases the type B agent's payoff in the Full-Censorship period since censoring is more costly, but increases her payoff in the Partial-Censorship period. Her total equilibrium payoff is the average payoff in those two periods. Therefore, if she is better off in equilibrium, it means the Full-Censorship period must be long and important. Hence, an increase in the cost decreases her equilibrium payoff. The inverse statement is also true. If the type B agent is worse off in equilibrium, it means the Full-Censorship period must be short and not important. In this case, an increase in the cost increases her equilibrium payoff. Therefore, the type B agent's payoff in equilibrium is monotonic in the cost, thus she prefers either a very strong institution or a very weak institution, depending on how long the Full-Censorship period is. The result is summarized below.

Corollary 3. *Assume $\beta^G = 0$, $c < \underline{c}$ and $\gamma < \beta^B$.*

- *Whether the type B agent is better off in the equilibrium than she is in the NCB does not depend on the censoring cost c;*
- *When $p_0 > p_{fast}$, the type B agent's expected payoff in equilibrium is monotonic in the censoring cost c. Moreover, it is decreasing (resp. increasing) in the cost, if and only if, she is better off (resp. worse off) in equilibrium than in the NCB.*

5 Inconclusive bad news

In this section, we consider the case where the good agent also has some bad news, but less frequently than the bad agent has (i.e. $0 < \beta^G < \beta^B$). Hence, in the NCB,[16] a piece of bad news is an informative but imperfect signal; the posterior public belief after a piece of bad news jumps down, but is above 0. Hence, after a piece of bad news, the evaluator knows that the agent is more likely to be a bad type, but he is not certain. The optimal policy for the evaluator's decision problem in the NCB is still a cutoff strategy, in which the cutoff belief depends on the information structure. However, since bad news is a noisy signal, the evaluator may make both type I and type II errors based on bad news.

Inconclusive bad news both enriches and complicates the problem. We will only focus on new insights derived from inconclusive bad news, but will not fully characterize all equilibria, nor conduct case-by-case discussion. In particular, we will show that it is possible that censorship may benefit the evaluator, as well as both types of the agent.

5.1 Faster good news

5.1.1 Equilibrium

We begin with the faster good news case; that is, good news arrives faster than the type B agent's bad news (i.e. $\gamma > \beta^B$). In the NCB, the public belief will drift down if no news arrives.[17]

Now, both types of agent need to choose a censorship strategy. The basic trade-off does not change – the agent censors bad news if and only if her equilibrium continuation value is higher than the censoring cost. The cost is the same for both types of agent, however, the continuation value is different. The good agent has a higher continuation value than the bad agent due to the following two reasons. First, since the good agent has a chance to succeed in her project, she will stay in her job longer in expectation than the bad agent. Second, the same censorship strategy is less costly for the good agent than the bad agent, since bad news happens more often to the bad agent. Therefore, it is the good agent who has a higher incentive to censor bad news. This observation significantly changes the welfare effect of censorship, which will be

[16] Though the NCB in the inconclusive news case is different from the NCB in the conclusive news case, we define them in the same way – the NCB means censorship is not possible, for example, due to a prohibitively high censoring cost.

[17] In the NCB, the direction of the belief drifting process is determined by the sign of $\gamma + \beta^G - \beta^B$. However, as we will show later, the corresponding equilibrium term is still determined by the sign of $\gamma - \beta^B$, since the type G agent censors bad news more aggressively than the type B agent does.

discussed after we introduce the following equilibrium.

Proposition 6. *Assume $c < \underline{c}$ and $\gamma > \beta^B$. There exists a $\bar{\beta} \in (0, \beta^B)$, such that for any $\beta^G \in (0, \bar{\beta}]$, there exists a PBE in which every player uses a pure cutoff strategy.*

1. *The evaluator dismisses the agent whenever the public belief is below p_{fast};*
2. *The type G agent censors bad news whenever the public belief is above p^G (but not equal to 1), where $p^G > p_{fast}$;*
3. *The type B agent censors bad news whenever the public belief is above $p^{B\dagger}$, where $p^{B\dagger} > p^G$.*

Here we maintain the assumption that the censoring cost is low (i.e. $c < \underline{c}$). In addition, we assume that the good agent's bad news does not happen too often (i.e. $\beta^G \in (0, \bar{\beta}]$). Figure 5 depicts the evolution of the public belief in the absence of news in equilibrium (blue dotted curve) and in the NCB (green solid curve). The first observation is that

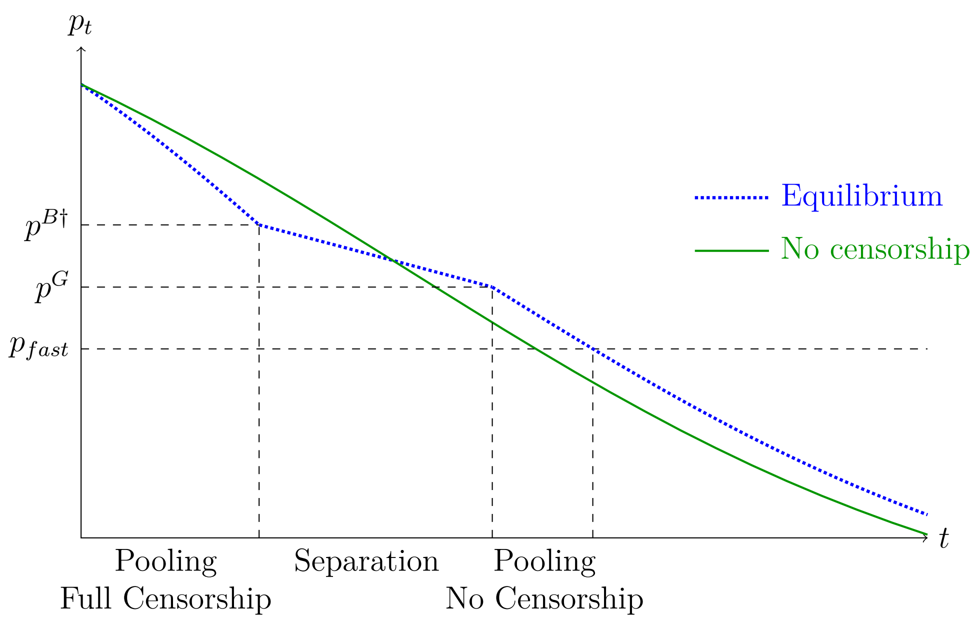

Figure 5: The belief evolution in the absence of news (inconclusive bad news & faster good news)

the evaluator's strategy is still a cutoff strategy with the cutoff belief p_{fast} both in the NCB and in equilibrium, due to the same reason in the conclusive bad news case, i.e. the cutoff belief does not depend on the arrival rate of bad news, thus does not depend on the censorship policy. Second, there are three different phases in equilibrium. Here, p^G (resp. $p^{B\dagger}$) is the cutoff belief below which the type G (resp. B) agent stops censoring. When the public belief is high (i.e. $p_t > p^{B\dagger}$), both types of agent censor bad news for sure – they are pooling on full censorship. Hence, there is no bad news in equilibrium. This implies that the belief drifts down faster than it does in the NCB (i.e. $\gamma > \gamma + \beta^G - \beta^B$). The off-path belief after a piece of bad news in this phase still needs to be specified. We set the public belief jump down to 0 after a piece of bad news in this phase. When the public belief is intermediate (i.e. $p_t \in (p^G, p^{B\dagger})$), only the type G agent censors for sure, but the type B agent stops censoring; we call it the Separation period. As we explained, since the type G agent has a higher incentive to censor bad news, she censors more aggressively than the type B agent in the Separation period. In equilibrium, bad news may arise in this phase. Since the separation between the two types of agent, a piece of bad news is endogenously conclusive. When it arrives, the public belief jumps down to 0 (i.e. $J(p_t, 1, 0) = 0$). The noisy inconclusive bad news becomes to a perfect conclusive signal because of the separation of censorship strategies. It also implies that the belief drifts down slower than it does in the NCB (i.e. $\gamma - \beta^B < \gamma + \beta^G - \beta^B$). This is due to the martingale property of public belief – the jumping process is amplified, thus the drifting process must be mitigated. When the public belief is low (i.e. $p_t < p^G$), both types of agent stops censoring – they are pooling on no censorship. In this phase, bad news is never censored by any type of agent. A piece of bad news makes the belief jump down from p_t to $j(p_t)$. We require $\beta^G \leq \bar{\beta}$, which is equivalent to $j(p^G) \leq p_{fast}$. Hence, in the last phase, a piece of bad news makes the public belief jump down below p_{fast}, which results in the dismissal of the agent.

5.1.2 Welfare: Equilibrium versus NCB

We now compare all players' expected payoffs from the above equilibrium and their expected payoffs in the NCB. We find that, under some conditions, every player has a strictly higher payoff from the above equilibrium than he or she has in the NCB.

Proposition 7. *Assume $c < \underline{c}$ and $\gamma > \beta^B$. There exists a $\underline{\beta} \in (0, \bar{\beta})$, when $\beta^G \in (0, \underline{\beta}]$ and $p_0 \in (p^G, p^{B\dagger}]$, the evaluator and both types of agent have strictly higher payoffs in the equilibrium of Proposition 6, compared with their payoffs in the NCB.*

The conditions we need for the above result are first the game starts from the Separation period (i.e. $p_0 \in (p^G, p^{B\dagger}]$), and second the good agent's bad news does not happen too often (i.e. $\beta^G \in (0, \underline{\beta}]$).[18] In the Separation period, due to the differential censorship strategies, a piece of bad news becomes a conclusive signal that the agent is type B. This improves the quality of the evaluator's information. In this period, he would never dismiss a type G agent, and the agent he might dismiss must be a type B agent. The endogenously conclusive signal helps the evaluator to avoid making type I and type II errors. This is why his payoff is strictly higher than his payoff in the

[18] Those are sufficient, but not necessary conditions to obtain the above result. By continuation, the above result still holds when the prior belief is slightly higher than $p^{B\dagger}$, but we restrict to those conditions to make a clear explanation about the underling economic forces.

NCB, if the game starts from the Separation period. Figure 6 shows the evaluator's value function in equilibrium, and the horizontal axis is the prior belief. The value functions in equilibrium and in the NCB may cross when the prior belief is high enough, but not in this numerical example. When the prior belief is high ($p_0 > p^{B\dagger}$), the evaluator expects to have low quality information first, i.e. no bad news can arise, then high quality information, i.e. conclusive bad news can arise. Whether he is better off in equilibrium compared with the NCB depends on whether the benefit from the Separation period dominates the loss from the period when both types of agent pool on full censorship.

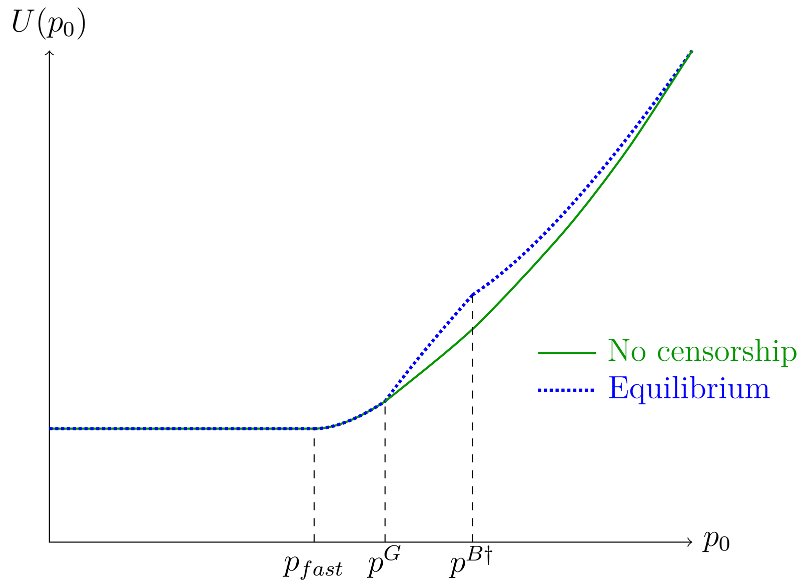

Figure 6: The evaluator's value functions (inconclusive bad news & faster good news)

Another implication from the Separation period is that no news becomes less severe, which is why both types of agent are better off in equilibrium. Though the public belief still drifts down in the absence of news, it drifts down slower than it does in the NCB. It means the evaluator would give more time to the agent in her job if no news arises, compared with the NCB. For the type B agent, when the game starts from the Separation period, she does not censor bad news both in the equilibrium and in the NCB. We also require $\beta^G \leq \underline{\beta}$, which is equivalent to $j(p^{B\dagger}) \leq p_{fast}$. Hence, a piece of bad news makes public belief jump down below p_{fast}, which results in the dismissal of the type B agent. Thus, the type B agent has the same payoff both in the equilibrium and in the NCB when a piece of bad news arrives. However, if the type B agent is lucky that bad news does not arise, she can stay longer in her job in the equilibrium, compared with the NCB. Thus, the type B agent is strictly better off in equilibrium. For the type G agent, she is also strictly better off in equilibrium. On the one hand side, similar to the conclusive bad news case, censoring bad news in the Separation period gives her a higher expected flow payoff in equilibrium. On the other hand side, she is given more time in her job if no news arrives in equilibrium, thus she is also more likely to succeed. Notice that censorship has opposite effect in the conclusive and the inconclusive bad news cases. In the conclusive news case, no news makes the belief drift down faster, while in the inconclusive news case, no news makes the belief drift down slower. Thus, the type G agent has a strictly higher payoff in the equilibrium than she has in the NCB.

When the censoring cost increases, the type B agent first gives up censorship, since she has a lower value in her job than the type G agent. If the censoring cost is intermediate (i.e. $c \in (\underline{c}, \bar{c})$, where $\bar{c} := \frac{\gamma+\rho_0}{\beta^G+\gamma+\rho_0}\frac{w}{\rho_0}$) and the type G agent's bad news does not happen too often (i.e. $\beta^G \leq \bar{\beta}$), we can show that there exists an equilibrium in which the evaluator and the type G agent use the same strategy as in Proposition 6, and the type B agent never censors. Notice that $\underline{c}$ is the cost threshold for the type B agent that determines whether she is willing to censor all bad news in exchange for staying in her job forever, while $\bar{c}$ is the cost threshold for the type G agent that determines whether she is willing to censor all bad news before she succeeds in exchange for staying in her job forever. Hence, the pooling on full censorship period disappears. We only have the Separation period and the pooling on no censorship period in equilibrium. Thus, we still have a similar welfare effect as in Proposition 7 under similar conditions.

5.2 Slower good news

5.2.1 Equilibrium

In this section, we consider the slower good news case, i.e. good news arrives slower than the type B agent's bad news ($\gamma < \beta^B$). "Slower good news" here does not necessary mean "no news is good news" in the NCB. The type G agent now has two possible signals – good news with intensity γ and bad news with intensity β^G, while the type B agent has only bad news with intensity β^B. Only when the signal intensity from the type G agent is smaller than the signal intensity from the type B agent (i.e. $\gamma + \beta^G < \beta^B$), "no news is good news", i.e. the public belief will drift up if no news arrives. Hence, when the good news arrives very slow (i.e. $\gamma < \beta^B - \beta^G$), "no news is good news" in the NCB. However, if good news arrives just mildly slow (i.e. $\beta^B - \beta^G < \gamma < \beta^B$), the public belief will drift down if no news arrives in the NCB. In both cases, the evaluator's optimal policy in the NCB is a cutoff strategy; he dismisses the agent whenever the public belief is below the cutoff $p^* \in (p_{slow}, p_{fast}]$. The cutoff belief p^* is strictly below p_{fast} when $\gamma < \beta^B - \beta^G$, and it is equal to p_{fast} when $\gamma \in (\beta^B - \beta^G, \beta^B)$.

Similar to the faster good news case, the type G agent has a higher incentive to censor bad news than the type B agent does. If we maintain the assumption that the censoring cost is low, $c < \underline{c}$, then the cutoff equilibrium in the conclusive bad news case still exists when bad news is actually inconclusive. That is, there exists an equilibrium in which the evaluator and the type B agent use the same strategies as in Proposition 4, and the type G agent censors bad news if and only if the public belief is above or equal to p_{fast}. First, it is easy to see that, given the strategies of both types of agent, the evaluator faces the exact same problem as in the conclusive bad news case when the public belief is above or equal to p_{fast}, since the type G agent censors all bad news. In addition, it is not hard to show that the evaluator wants to dismiss the agent whenever the public belief is below p_{fast}. Second, given the evaluator's strategy, the type B agent faces the exact same problem as in the conclusive bad news case, hence she also uses the same strategy. At last, since the type G agent has a higher incentive to censor bad news than the type B agent does, the former censors bad news with probability one whenever the latter censors with positive probability. Therefore, the equilibrium in the conclusive bad news case still exists.

When the censoring cost increases, the type B agent first gives up censorship. When the cost is intermediate (i.e. $c \in (\underline{c}, \bar{c})$), there exists an equilibrium in which the type B agent never censors, but the type G agent censors bad news whenever the public belief is above p_{slow}. In addition, the equilibrium strategy of the evaluator is to dismiss the agent whenever the public belief is below p_{slow}. It is worth emphasizing that p_{slow} is the cutoff belief in the evaluator's optimal policy in the NCB when bad news is conclusive. In this equilibrium, separation between the two types of agent makes a piece of bad news endogenously conclusive. Hence, the evaluator's equilibrium strategy is the cutoff strategy with the cutoff belief p_{slow}. In this equilibrium, since good news arrives slower than the type B agent's bad news, "no news is good news"; the public belief drifts up in the absence of news. Since the censoring cost is intermediate, it is too costly for the type B agent to censor any news, but it is worth for the type G agent censoring bad news in exchange for staying in the job.

The following proposition summarizes the result.

Proposition 8. *Assume $\gamma < \beta^B$ and $\beta^G \in (0, \beta^B)$.*

1. *(Low cost) When $c < \underline{c}$, there exists an equilibrium as follows:*
   - *The evaluator and the type B agent use the same strategies as in Proposition 4;*
   - *The type G agent censors bad news whenever the public belief is above and equal to p_{fast}, but below 1.*
2. *(Intermediate cost) When $c \in (\underline{c}, \bar{c})$, there exists an equilibrium as follows:*
   - *The evaluator dismisses the agent whenever the public belief is below or equals to p_{slow};*
   - *The type G agent censors bad news whenever the public belief is above p_{slow}, but below 1, and the type B agent never censors.*

5.2.2 Welfare: Equilibrium versus NCB

We still want to compare equilibrium payoffs with the payoffs in the NCB. However, the comparison would be tedious if we try to cover every possible scenario, especially now the NCB has two very different cases, depending on whether the good news arrives very slow or mildly slow. Hence, we only give a big picture to see how the change of the censoring cost may change the welfare comparison from the evaluator's point of view.

Proposition 9. *Assume $\gamma < \beta^B$ and $\beta^G \in (0, \beta^B)$.*

1. *(Low cost) Assume $c < \underline{c}$. When $p_0 \in (p^*, 1)$, the evaluator has a strictly lower payoff in the equilibrium from Proposition 8 than he has in the NCB.*

2. *(Intermediate cost) Assume $c \in (\underline{c}, \bar{c})$. When $p_0 \in (p_{slow}, 1)$, the evaluator has a strictly higher payoff in the equilibrium from Proposition 8 than he has in the NCB.*

Figure 7 illustrates the evaluator's value functions in the NCB and in the equilibrium from Proposition 8. When the censoring cost is low, since both types of agent censors very aggressively, the evaluator essentially cannot rely on bad news to learn the type of the agent. This is the worst information he could have, thus he is always worse off in the equilibrium, compared with the NCB. In addition, the evaluator's incentive for learning is discouraged when the good news arrives very slow. That is, when $\gamma < \beta^B - \beta^G$, the cutoff belief in the evaluator's optimal policy in the NCB increases from p^* to the cutoff belief in his equilibrium strategy p_{fast}. We have seen this discouragement effect from the conclusive bad news case. When the censoring cost is intermediate, the differential censorship policy between two types of agent dramatically improves the quality of information – a piece of bad news is endogenously conclusive. Hence, the evaluator is better off in the equilibrium, compared with the NCB. Moreover, the improvement in the quality of information encourages the evaluator to explore the type of the agent even when the public belief is in between p_{slow} and p^*. We call it the *encouragement effect*; the cutoff belief in the evaluator's optimal policy in the NCB decreases from p^* to the cutoff belief in his equilibrium strategy p_{slow}.

Both types of agent are also directly affected by the discouragement and the encouragement effects of the evaluator. When the censoring cost is intermediate, because of the encouragement effect, both types of agent are better off in the equilibrium when the prior belief is in between p_{slow} and p^*. In the equilibrium, the type G agent will stay in her job forever, and the type B agent will stay in her job for some positive time. While in the NCB, both types of agent will not have a chance to start. When the censoring cost is low and the good news arrives very low ($\gamma < \beta^B - \beta^G$), the effect is opposite; both types of agent are worse off in the equilibrium when the prior belief is in between p^* and p_{fast}, due to the discouragement effect.

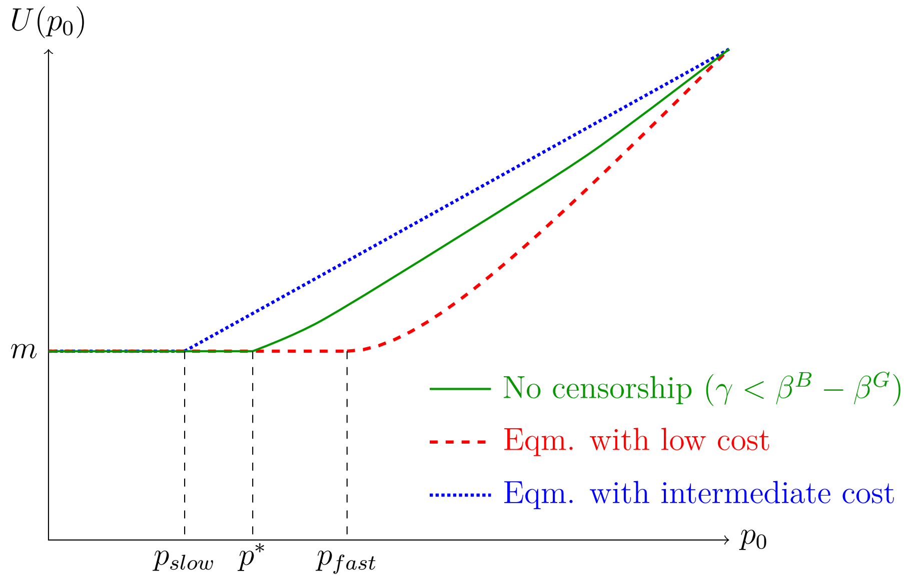

Figure 7: The evaluator's value functions (inconclusive bad news & slower good news)

6 Conclusion

This paper studies costly censorship in a dynamic environment. An evaluator tries to learn the agent's competence from a good news process and a bad news process. However, an informed agent who knows her competence level can conceal bad news at some cost.

When bad news is conclusive, only the incompetent agent censors. The evaluator suffers from censorship since otherwise valuable information has been suppressed. This also makes her interpret no news more severely, which hurts the competent agent. It may also hurt the incompetent agent, even though censorship helps her survive bad news. The incompetent agent is worse off with censorship when the period that she finds it optimal to censor is short. When good news arrives slower than the bad news from the incompetent agent, a discouragement effect occurs – the evaluator increases the threshold public belief of dismissal.

When bad news is inconclusive, both competent and incompetent agents censor bad news. However, it is the competent agent who has a higher incentive to censor since she has a higher value in the job. When the two types of agent separate in their censorship strategies (i.e. only the competent agent censors), the quality of information for the evaluator is improved – the inconclusive noisy bad news becomes conclusive and perfect. Not only the evaluator benefits from that, but also both types of agent, since no news now is less severe and the agent may stay in her job longer. When good news arrives very slow, an encouragement effect may occur – the evaluator decreases the threshold public belief of dismissal, when the censoring cost is not too high or too low such that only the competent agent finds it optimal to censor.

Appendices

Proof of Proposition 1. (1) First, we show that when $c > \underline{c}$, the type B agent never censors bad news under any history $\bar{\varnothing}_t^B$ in any PBE.

In any PBE, a piece of bad news results in an immediate dismissal of the agent. Fixed any equilibrium and the equilibrium strategies $\{\hat{\boldsymbol{r}}, \hat{\boldsymbol{x}}^B\}$, let $V_B^{\hat{\boldsymbol{r}}}(\bar{\varnothing}_t^B) \geq 0$ be the equilibrium payoff of the type B agent under the history $\bar{\varnothing}_t^B$. Suppose $\hat{x}_t^B > 0$ for a history $\bar{\varnothing}_t^B$, then under alternative strategy of the evaluator $\boldsymbol{r}_\infty$, the type B agent obtains $v_B^{\boldsymbol{r}_\infty, \hat{\boldsymbol{x}}^B}(\bar{\varnothing}_t^B)$ by censoring that piece of arrived bad news. On the other hand, under the evaluator's strategy $\boldsymbol{r}_\infty$, if the agent censors that piece of arrived bad news and resumes the optimal strategy (i.e. never censors again), she can obtain

$$-\rho_0 c + \mathbb{E}[\int_0^{\tau^B} \rho_0 e^{-\rho_0 \nu} w \, \mathrm{d}\nu + e^{-\rho_0 \tau^B} V_B^{\boldsymbol{r}_\infty}(\bar{\varnothing}_{\tau^B}^B)] = \frac{\rho_0}{\rho_0 + \beta^B}[w - (\rho_0 + \beta)c],$$

where τ^B is the arrival time of the next piece of bad news from the type B agent. Thus, $v_B^{\boldsymbol{r}_\infty, \hat{\boldsymbol{x}}^B}(\bar{\varnothing}_t^B) \leq \frac{\rho_0}{\rho_0 + \beta^B}[w - (\rho_0 + \beta)c] < 0$, since $c > \underline{c}$.

In addition, $V_B^{\hat{\boldsymbol{r}}}(\bar{\varnothing}_t^B) \leq v_B^{\boldsymbol{r}_\infty, \hat{\boldsymbol{x}}^B}(\bar{\varnothing}_t^B) < 0$ since any strategy of the evaluator is weakly worse for the agent than $\boldsymbol{r}_\infty$ as the latter subscribes that the evaluator never dismisses the agent without bad news. This contradicts $V_B^{\hat{\boldsymbol{r}}}(\bar{\varnothing}_t^B) \geq 0$. Thus, $\hat{x}_t^B = 0$ for any history $\bar{\varnothing}_t^B$.

(2) Second, since the type B agent never censors bad news, the evaluator faces a standard exponential bandit model. The literature on bandit models provides the optimal strategy of the evaluator. For completeness, we sketch the proof here.

Let $\mathbf{0}$ represent the strategy of the type B agent that never censors bad news, i.e. $x_t^B = 0$ for any t. Then the evaluator chooses $\boldsymbol{r}_t$ to maximize

$$\begin{aligned} u^{\boldsymbol{r}, \mathbf{0}}(\varnothing_t) &= \mathbb{E}\left[\int_t^T \rho_1 e^{-\rho_1(\nu - t)} \, \mathbb{1}_{\theta = G} \, k \, \mathrm{d}N_\nu^G + e^{-\rho_1(T-t)} m \,\middle|\, \varnothing_t\right] \\ &= \mathbb{E}\left[\int_t^T \rho_1 e^{-\rho_1(\nu - t)} p_\nu h \, \mathrm{d}\nu + e^{-\rho_1(T-t)} m \,\middle|\, \varnothing_t\right], \end{aligned}$$

where the second equation comes from the Law of Iterated Expectations.

Since the payoff of the evaluator depends on history only through the public belief p_t, we rewrite the above $u^{\boldsymbol{r}, \mathbf{0}}(p_t) := u^{\boldsymbol{r}, \mathbf{0}}(\varnothing_t)$, and

$$U^{\mathbf{0}}(p_t) = \sup_{\boldsymbol{r}_t} u^{\boldsymbol{r}, \mathbf{0}}(p_t).$$

When $\gamma > \beta^B$, the following ODE solves $U^{\mathbf{0}}(p)$ when it is optimal to experiment,

$$\rho_1 U^{\mathbf{0}}(p) = \rho_1 ph + p\gamma(h - U^{\mathbf{0}}(p)) + (1-p)\beta^B(m - U^{\mathbf{0}}(p)) - p(1-p)(\gamma - \beta^B)U^{\mathbf{0}\prime}(p).$$

When the evaluator is indifferent between experimenting and dismissal at a public belief $\hat{p}$, then $U^{\mathbf{0}}(\hat{p}) = m$ (value matching) and $U^{\mathbf{0}\prime}(\hat{p}) = 0$ (smooth pasting) since the continuation region and stopping region communicate. Hence,

$$\rho_1 m = \rho_1 \hat{p} h + \hat{p}\gamma(h - m),$$

i.e. $\hat{p} = p_{fast} = \frac{\rho_1 m}{\rho_1 h + \gamma(h-m)}$. In addition, we can show that when $p > p_{fast}$,

$$U^{\mathbf{0}}(p) = m + p(h-m)(1 - e^{-(\rho_1+\gamma)\bar{s}}) - (1-p)m\frac{\rho_1}{\rho_1 + \beta^B}(1 - e^{-(\rho_1+\beta^B)\bar{s}}).$$

Here, $\bar{s} = \frac{\ln[\frac{p}{1-p}\frac{1-p_{fast}}{p_{fast}}]}{\gamma - \beta^B}$ is the duration when the evaluator experiments without a conclusive signal.

When $\gamma < \beta^B$, the same ODE solves $U^{\mathbf{0}}(p)$ when it is optimal to experiment. However, the smooth pasting condition does not hold, since the continuation region and stopping region do not communicate. A more direct approach is to compute the optimal experimenting time without a conclusive signal. If the experimenting time is s, then the expected payoff of the evaluator is

$$m + p(h-m)(1 - e^{-(\rho_1+\gamma)s}) - (1-p)m\frac{\rho_1}{\rho_1 + \beta^B}(1 - e^{-(\rho_1+\beta^B)s}).$$

It is easy to see (by taking first order derivative with respect to s) that it is either optimal to experiment forever or to stop immediately. Therefore, value matching gives $\hat{p} = p_{slow} = \frac{\rho_1 m}{\rho_1 h + \beta^B(h-m)}$ which solves

$$m + \hat{p}(h-m) - (1-\hat{p})m\frac{\rho_1}{\rho_1 + \beta^B} = m.$$

In addition, when $p > p_{slow}$,

$$U^{\mathbf{0}}(p) = m + p(h-m) - (1-p)m\frac{\rho_1}{\rho_1 + \beta^B}.$$

□

Proof of Proposition 2. (1) First, for any admissible strategy of the type B agent, the evaluator's best response, if it exists, is a pure cutoff strategy with the cutoff p_{fast}.

Let **0** and **1** represent the strategy of the type B agent that never censors bad news ($x_t^B = 0$ for any t) and always censors bad news ($x_t^B = 1$ for any t). Standard analysis in bandit literature shows that in both cases the optimal response of the evaluator is to use a cutoff strategy with the cutoff $p_{fast} = \frac{\rho_1 m}{\rho_1 h + \gamma(h-m)}$. We have seen the case without censorship. To see the case with full censorship, let us write $U^{\mathbf{1}}(p_0)$ as the evaluator's value function. When it is optimal to experiment, it solves the following ODE,

$$\rho_1 U^{\mathbf{1}}(p) = \rho_1 ph + p\gamma(h - U^{\mathbf{1}}(p)) - p(1-p)\gamma U^{\mathbf{1}\prime}(p).$$

Value matching and smooth pasting give $\hat{p} = p_{fast}$ which solves

$$\rho_1 m = \rho_1 ph + p\gamma(h-m).$$

In addition, when $p > p_{fast}$

$$U^{\mathbf{1}}(p) = m + p(h-m)(1 - e^{-(\rho_1+\gamma)\hat{s}}) - (1-p)m(1 - e^{-\rho_1 \hat{s}}).$$

Here, $\hat{s} = \frac{\ln[\frac{p}{1-p}\frac{1-p_{fast}}{p_{fast}}]}{\gamma}$ is the duration when the evaluator experiments without a conclusive signal.

There two cases represent the maximal and minimal information that the evaluator can get. The minimal information is a garbling of the information that generated from any other admissible strategy of the type B agent, which in turn is a garbling of the maximal information.

Consider the case where the type B agent uses a strategy $\boldsymbol{x}^B$, and the public conjecture is $\tilde{\boldsymbol{x}}^B$. Let $\boldsymbol{r}$ be the best response of the evaluator with respect to $\tilde{\boldsymbol{x}}^B$. In addition, let $\boldsymbol{r}_0$ and $\boldsymbol{r}_1$ be the best responses of the evaluator with respect to $\mathbf{0}$ and $\mathbf{1}$, respectively. Then, we have

$$U^{\mathbf{0}}(p) = u^{\boldsymbol{r}_0,\mathbf{0}}(p) \geq u^{\boldsymbol{r},\mathbf{0}}(p) \geq U^{\tilde{\boldsymbol{x}}^B}(p) = u^{\boldsymbol{r},\tilde{\boldsymbol{x}}^B}(p) \geq u^{\boldsymbol{r}_1,\tilde{\boldsymbol{x}}^B}(p) \geq u^{\boldsymbol{r}_1,\mathbf{1}}(p) = U^{\mathbf{1}}(p).$$

This is because by using the same strategy, the evaluator has weakly higher value under the maximal information than under any other information induced by the strategy of the type B agent, and he has weakly lower value under the minimal information than under any other information.

In addition, when $p > p_{fast}$, the best response is to keep experimenting since it is also the best response under the minimal information. Similarly, the best response is to dismiss the agent when $p < p_{fast}$.

(2) Second, given the evaluator uses a cutoff strategy with the cutoff p_{fast} in any PBE, the type B agent would be eventually dismissed. Let $s_1 = \sup\{t \geq 0 : \bar{h}_t^B \in \bar{\varnothing}_t^B, x_t^B > 0\}$ in a PBE. Clearly, without any revealed signal until s_1, we must have $p_{s_1} > p_{fast}$ in any PBE, otherwise censoring a piece of bad news only incurs a cost of c, but gives 0 continuation value to the agent. Hence, when a piece of bad news arrives at time $t > s_1$, $x_t^B = 0$.

In addition, when bad news has not been revealed to the evaluator yet at time $t < s_1$, we must have $x_t^B = 1$. Suppose not, i.e. there exists a $t < s_1$ and $x_t^B < 1$. Since $s_1 =$

$\sup\{t \geq 0 : \bar{h}_t^B \in \bar{\varnothing}_t^B, x_t^B > 0\}$, there must a $t' \in (t, s_1]$ such that $x_t' > 0$. Obviously, the agent's equilibrium payoff cannot be negative, $V_B^{\boldsymbol{r}}(\bar{\varnothing}_{t'}^B) \geq 0$. Equivalently,

$$\mathbb{E}\left[\int_{t'}^{T} \rho_0 e^{-\rho_0(\nu - t')}(w \,\mathrm{d}\nu - cX_\nu^B \,\mathrm{d}N_\nu^B)\bigg|\, \bar{\varnothing}_{t'}^B\right] \geq \rho_0 c,$$

where T is induced by the equilibrium strategy of the evaluator $\boldsymbol{r}$. Consider an alternative strategy $\boldsymbol{x}^{B'}$ such that censor all bad news for time $\nu \in [t, t']$ and resume the original strategy, then

$$\begin{aligned}
v_B^{\boldsymbol{r},\boldsymbol{x}^{B'}}(\bar{\varnothing}_t^B) &= -\rho_0 c + \int_t^{t'} \rho_0 e^{-\rho_0(\nu - t)}(w \,\mathrm{d}\nu - \beta^B c \,\mathrm{d}\nu) \\
&\quad + e^{-\rho_0(t'-t)}\,\mathbb{E}\left[\int_{t'}^{T} \rho_0 e^{-\rho_0(\nu - t')}(w \,\mathrm{d}\nu - cX_\nu^B \,\mathrm{d}N_\nu^B)\bigg|\, \bar{\varnothing}_{t'}^B\right] \\
&\geq -\rho_0 c + (w - \beta^B c)(1 - e^{-\rho_0(t'-t)}) + e^{-\rho_0(t'-t)}\rho_0 c \\
&= [w - (\rho_0 + \beta^B)c](1 - e^{-\rho_0(t'-t)}) > 0.
\end{aligned}$$

Hence, it is strictly optimal to censor that piece of bad news, which contradicts with $x_t^B < 1$.

At time s_1, note that

$$V_B^{\boldsymbol{r}}(\bar{\varnothing}_{s_1}^B) = \max\left\{\mathbb{E}[-\rho_0 c + \int_0^{s_2 \wedge \tau^B} \rho_0 e^{-\rho_0 \nu} w \,\mathrm{d}\nu], 0\right\},$$

where $s_2 = \frac{\ln[\frac{p_{s_1}}{1-p_{s_1}}\frac{1-p_{fast}}{p_{fast}}]}{\gamma - \beta^B}$ is duration when the public belief drifts from p_{s_1} to p_{fast} in the absence of news, and τ^B is the arrival time of the next piece of bad news. Clearly, the first term is continuous and increasing in s_2. Since for any $s < s_2$, $\mathbb{E}[-\rho_0 c + \int_0^{s \wedge \tau^B} \rho_0 e^{-\rho_0 \nu} w \,\mathrm{d}\nu] \leq 0$, we must have

$$\mathbb{E}[-\rho_0 c + \int_0^{s_2 \wedge \tau^B} \rho_0 e^{-\rho_0 \nu} w \,\mathrm{d}\nu] \leq 0.$$

Suppose that the above expression is strictly less than 0; let is be $-\epsilon < 0$. Then consider the value at time $t < s_1$. Since it is optimal to censor any news before s_1 and censors no news after s_1, thus

$$\begin{aligned}
V_B^{\boldsymbol{r}}(\bar{\varnothing}_t^B) &= -\rho_0 c + \int_0^{s_1 - t} \rho_0 e^{-\rho_0 \nu}(w - \beta^B c)\,\mathrm{d}\nu + e^{-\rho_0(s_1 - t)} \int_0^{s_2 \wedge \tau^B} \rho_0 e^{-\rho_0 \nu} w \,\mathrm{d}\nu \\
&= -\rho_0 c + (w - \beta^B c)(1 - e^{-\rho_0(s_1 - t)}) + e^{-\rho_0(s_1 - t)}(\rho_0 c - \epsilon) \\
&= [w - (\rho_0 + \beta^B)c] - e^{-\rho_0(s_1 - t)}[w - (\rho_0 + \beta^B)c + \epsilon].
\end{aligned}$$

Clearly, when t is sufficiently close to s_1, the above expression is less than 0, which is a contradiction. Hence, we must have

$$\mathbb{E}[-\rho_0 c + \int_0^{s_2 \wedge \tau^B} \rho_0 e^{-\rho_0 \nu} w \,\mathrm{d}\nu] = 0.$$

This pins down

$$s_2 = \frac{\ln \frac{w}{w-(\rho_0+\beta^B)c}}{\rho_0 + \beta^B},$$

which in turn determines

$$p^B = p_{s_1} = \frac{p_{fast}}{p_{fast} + (1-p_{fast})e^{-(\gamma-\beta^B)s_2}}.$$

Hence, if $p_0 > p^B$, the type B ruler would censor all bad news for s_1 time period,

$$s_1 = \frac{\ln[\frac{p_0}{1-p_0}\frac{1-p^B}{p^B}]}{\gamma} = (1-\frac{\beta^B}{\gamma})(\bar{s}-s_2).$$

If $p_0 \leq p^B$, the type B ruler would not censor any bad news.

Hence, the type B ruler's best response is also a cutoff strategy with the cutoff p^B. It is easy to verify that the best response of the evaluator to such a cutoff strategy of the type B ruler exists, and it is precisely the cutoff strategy with the cutoff p_{fast}.

Since the strategy of the type B ruler is also unique in any PBE, this is the unique PBE. $\square$

Proof of Proposition 3. When $p_0 > p^B$, we can compute the expected payoff of the type B ruler in the equilibrium,

$$(w-\beta^B c)(1-e^{-\rho_0 s_1})+e^{-\rho_0 s_1}\frac{\rho_0 w}{\rho_0+\beta^B}(1-e^{-(\rho_0+\beta^B)s_2}) = (w-\beta^B c)(1-e^{-\rho_0 s_1})+e^{-\rho_0 s_1}\rho_0 c.$$

If the censorship is not possible, then her expected payoff would be

$$\frac{\rho_0 w}{\rho_0+\beta^B}(1-e^{-(\rho_0+\beta^B)\bar{s}}).$$

Hence, she would be better off in the equilibrium with censorship if and only if

$$(w-\beta^B c)(1-e^{-\rho_0 s_1})+e^{-\rho_0 s_1}\rho_0 c > \frac{\rho_0 w}{\rho_0+\beta^B}(1-e^{-(\rho_0+\beta^B)\bar{s}}).$$

Since we have

$$w-\beta^B c > \frac{\rho_0 w}{\rho_0+\beta^B} > \frac{\rho_0 w}{\rho_0+\beta^B}(1-e^{-(\rho_0+\beta^B)\bar{s}}) > \rho_0 c,$$

thus she would be better off in the equilibrium with censorship if and only if

$$e^{-\rho_0 s_1} < \frac{\beta^B + \rho_0 e^{-\frac{\gamma(\rho_0+\beta^B)}{\gamma-\beta^B}s_1}}{\rho_0+\beta^B}.$$

Let

$$f(s) = e^{-\rho_0 s} - \frac{\beta^B + \rho_0 e^{-\frac{\gamma(\rho_0+\beta^B)}{\gamma-\beta^B}s}}{\rho_0 + \beta^B}.$$

Note that $f(s)$ is first increasing in s, then decreasing in s. Moreover, $f(0) = 0$, and $\lim_{s\to\infty} f(s) = -\frac{\beta^B}{\rho_0+\beta^B} < 0$. Hence, there exists a $s_1^* = s_1^*(\rho_0, \gamma, \beta^B) > 0$, such that the type B ruler would be worse off in the equilibrium if and only if $s_1 < s_1^*$. □

Proof of Corollary 1. Let $s^* := \frac{\gamma}{\gamma-\beta^B} s_1^*$. Since $s_1 = \frac{\gamma-\beta^B}{\gamma}(\bar{s} - s_2) > 0$, $s_1 < s_1^*$ is equivalent to $\bar{s} < s_2 + s^*$.

Note that

$$p^B = \frac{p_{fast}}{p_{fast} + (1 - p_{fast})e^{-(\gamma-\beta^B)s_2}}$$

is increasing in s_2. In addition, s_2 is increasing in c, and $\lim_{c\to 0} p^B = p_{fast}$ and $\lim_{c\to\bar{c}} p^B = 1$. Hence, $p_0 = p^B$ defines $\mathfrak{c}_3 \in (0, \underline{c})$, such that $p_0 \le p^B$ when $c \in [\mathfrak{c}_3, \underline{c})$. Thus, when $c \in [\mathfrak{c}_3, \underline{c})$, the equilibrium payoff and the payoff in the NCB would be the same for the type B agent.

The type B agent would be worse off in the equilibrium if and only if $s_2 > \bar{s} - s^*$. If $\bar{s} \le s^*$, then $s_2 > \bar{s} - s^*$ holds for all $c \in (0, \mathfrak{c}_3]$, and the type B agent would be worse off in the equilibrium. Suppose $\bar{s} > s^*$. Since s_2 is increasing in c, and $\lim_{c\to 0} s_2 = 0$ and $\lim_{c\to\mathfrak{c}_3} s_2 = \bar{s}$, $s_2 = \bar{s} - s^*$ defines $\mathfrak{c}_1 \in (0, \mathfrak{c}_3)$ such that $s_2 > \bar{s} - s^*$ when $c \in (\mathfrak{c}_1, \mathfrak{c}_3]$, and $s_2 < \bar{s} - s^*$ when $c \in (0, \mathfrak{c}_1)$. We complete the proof by redefining $\mathfrak{c}_1$ as $\mathfrak{c}_1 \mathbb{1}_{\bar{s}>s^*}$. □

Proof of Corollary 2. From Corollary 1, when $c \in [\mathfrak{c}_3, \bar{c})$, the equilibrium payoff and the NCB payoff would be the same for the type B ruler, and they do not depend on the cost c.

Suppose $c < \mathfrak{c}_3$, then $p_0 > p^B$. The expected payoff of the type B ruler in the equilibrium is

$$(w - \beta^B c)(1 - e^{-\rho_0 s_1}) + e^{-\rho_0 s_1}\rho_0 c.$$

Note that $s_1 = \frac{\gamma-\beta^B}{\gamma}(\bar{s} - s_2)$ and $\frac{\partial s_2}{\partial c} = \frac{1}{w-(\rho_0+\beta^B)c}$. Thus,

$$\frac{\partial(w - \beta^B c)(1 - e^{-\rho_0 s_1}) + e^{-\rho_0 s_1}\rho_0 c}{\partial c} = -\beta^B + \beta^B e^{-\rho_0 s_1}\frac{\rho_0 + \gamma}{\gamma}.$$

Hence, the type B ruler's equilibrium payoff is increasing in c if and only if

$$e^{-\rho_0 s_1} > \frac{\gamma}{\rho_0 + \gamma}.$$

Let $s_1^{**} = s_1^{**}(\rho_0, \gamma) > 0$ be the solution of $e^{-\rho_0 s} = \frac{\gamma}{\rho_0+\gamma}$. Hence, the type B ruler's equilibrium payoff is increasing (resp. decreasing) in c if and only if $s_1 < s_1^{**}$ (resp. $s_1 > s_1^{**}$).

Let $s^{**} = \frac{\gamma}{\gamma-\beta^B} s_1^{**} > 0$. Hence, $s_1 < s_1^{**}$ is equivalent to $s_2 > \bar{s} - s^{**}$. Note that s_2 is increasing in c, and $\lim_{c\to 0} s_2 = 0$ and $\lim_{c\to \mathfrak{c}_3} s_2 = \bar{s}$.

If $\bar{s} \leq s^{**}$, then $s_2 > \bar{s} - s^{**}$ holds for all $c \in (0, \mathfrak{c}_3]$, and the type B ruler's equilibrium payoff is increasing in $c \in (0, \mathfrak{c}_3]$. Suppose $\bar{s} > s^{**}$, then $s_2 = \bar{s} - s^{**}$ defines $\mathfrak{c}_2 \in (0, \mathfrak{c}_3)$ such that $s_2 > \bar{s} - s^{**}$ when $c \in (\mathfrak{c}_2, \mathfrak{c}_3]$, and $s_2 < \bar{s} - s^{**}$ when $c \in (0, \mathfrak{c}_2)$. Last, we redefine $\mathfrak{c}_2$ as $\mathfrak{c}_2 \, \mathbb{1}_{\bar{s} > s^{**}}$.

Finally, we show that $s^{**} \in (0, s^*)$, or equivalently $s_1^{**} \in (0, s_1^*)$. Let $\nu_1 = s_1^{**}$ be the solution of

$$e^{-\rho_0 \nu_1} = \frac{\gamma}{\rho_0 + \gamma},$$

and ν_2 be the solution of

$$\frac{\beta^B + \rho_0 e^{-\frac{\gamma(\rho_0+\beta^B)}{\gamma-\beta^B}\nu_2}}{\rho_0 + \beta^B} = \frac{\gamma}{\rho_0 + \gamma}.$$

Hence, $s_1^{**} < s_1^*$ is equivalent to $\nu_2 < \nu_1$. Note that $\nu_1 = \frac{\ln[\frac{\rho_0+\gamma}{\gamma}]}{\rho_0}$, and $\nu_2 = \ln[\frac{\rho_0+\gamma}{\gamma-\beta^B}]\frac{\gamma-\beta^B}{\gamma(\rho_0+\beta^B)}$. In addition,

$$\frac{\partial \nu_2}{\partial \beta^B} = \frac{1}{\gamma}\frac{1}{\rho_0 + \beta^B}(1 - \ln[\frac{\rho_0+\gamma}{\gamma-\beta^B}]\frac{\rho_0+\gamma}{\rho_0+\beta^B})$$

Let $\eta = \frac{\gamma-\beta^B}{\rho_0+\gamma} \in (0,1)$, then $\frac{\partial \nu_2}{\partial \beta^B} < 0$ if and only if

$$1 - \eta + \ln \eta < 0,$$

which is always true. Hence, ν_2 is decreasing in $\beta^B \in (0, \gamma)$. Also,

$$\lim_{\beta^B \to 0} \nu_2 = \nu_1.$$

Hence, for $\beta^B \in (0, \gamma)$, $\nu_2 < \nu_1$. Thus, $s_1^{**} < s_1^*$ and $s^{**} < s^*$. In addition, by the definitions of $\mathfrak{c}_1$ and $\mathfrak{c}_2$, $\mathfrak{c}_1 = \mathfrak{c}_2 = 0$ when $\bar{s} \leq s^{**}$, $\mathfrak{c}_1 = 0 < \mathfrak{c}_2$ when $\bar{s} < s^{**}$ when $\bar{s} \in (s^{**}, s^*]$, and $0 < \mathfrak{c}_1 < \mathfrak{c}_2$ when $\bar{s} > s^*$.

□

Proof of Proposition 4. (1) First, let us verify that the pair of mixed cutoff strategies constitutes an equilibrium. Given the type B agent's strategy, call it $\boldsymbol{x}_M$,

$$x_t^B = \begin{cases} 1 & \text{if } p_t > p_{fast}, \\ 1 - \frac{\gamma}{\beta^B} & \text{if } p_t = p_{fast}, \\ 0 & \text{if } p_t < p_{fast}, \end{cases}$$

we now solve the evaluator's best response. Consider three cases: $p_0 = p_{fast}$, $p_0 > p_{fast}$ and $p_0 < p_{fast}$.

When $p_0 = p_{fast}$, the public belief does not change in the absence of news. The evaluator faces a stationary problem, and his value function solves

$$U^{\boldsymbol{x}_M}(p_{fast}) = \max\left\{p_{fast}h + \frac{\gamma}{\rho_1}[p_{fast}h + (1-p_{fast})m - U^M(p_{fast})], m\right\}.$$

It is easy to check that

$$U^{\boldsymbol{x}_M}(p_{fast}) = \max\left\{p_{fast}h + (1-p_{fast})\frac{\gamma}{\rho_1+\gamma}m, m\right\}.$$

Note that the first term is exactly m, thus the evaluator is indifferent between experimenting or not.

When $p_0 > p_{fast}$, the public belief drifts down until p_{fast} in the absence of news. The strategy of the evaluator is just how long to experiment in the absence of any signal; we use $\boldsymbol{r}_s$ to denote this strategy, T_s to denote the stopping time, and $s \in [0, \hat{s}]$ to be the length of experimentation. Then his payoff would be

$$u^{\boldsymbol{r}_s,\boldsymbol{x}_M}(p_0) = \mathbb{E}\left[\int_0^{T_s} \rho_1 e^{-\rho_1 \nu} p_\nu h \,\mathrm{d}\nu + e^{-\rho_1 T_s} m\right].$$

We can show that

$$u^{\boldsymbol{r}_s,\boldsymbol{x}_M}(p_0) = p_0[\int_0^s \gamma e^{-\gamma\nu}\rho_1 e^{-\rho_1\nu}(k + \frac{h}{\rho_1})\,\mathrm{d}\nu + e^{-\rho_1 s}e^{-\gamma s}m] + (1-p_0)e^{-\rho_1 s}m.$$

Hence,

$$\begin{aligned}
\frac{\partial u^{\boldsymbol{r}_s,\boldsymbol{x}_M}(p_0)}{\partial s} &= p_0 e^{-(\rho_1+\gamma)s}(\rho_1+\gamma)(h-m) - (1-p_0)e^{-\rho_1 s}\rho_1 m \\
&= (1-p_0)e^{-\rho_1 s}[\frac{p_0}{1-p_0}e^{-\gamma s}(\rho_1+\gamma)(h-m) - \rho_1 m] \\
&= (1-p_0)e^{-\rho_1 s}[\frac{p_s}{1-p_s}(\rho_1+\gamma)(h-m) - \rho_1 m].
\end{aligned}$$

Since $p_s \in [p_{fast}, p_0]$, and $\frac{p_{fast}}{1-p_{fast}} = \frac{\rho_1 m}{(\rho_1+\gamma)(h-m)}$, we have $\frac{\partial u^{\boldsymbol{r}_s,\boldsymbol{x}_M}(p_0)}{\partial s} > 0$ for $s \in [0, \hat{s})$. Hence, it is optimal for the evaluator to experiment $\hat{s}$ time until the public belief drifts down to p_{fast}.

When $p_0 < p_{fast}$, the public belief drifts up until p_{fast} in the absence of news. Let it be $s \in [0, \bar{s}]$ be the length that the evaluator experiments without a conclusive signal. Then his payoff would be

$$\begin{aligned}
u^{\boldsymbol{r}_s,\boldsymbol{x}_M}(p_0) = p_0[&\int_0^s \gamma e^{-\gamma\nu}\rho_1 e^{-\rho_1\nu}(k + \frac{h}{\rho_1})\,\mathrm{d}\nu + e^{-\rho_1 s}e^{-\gamma s}m] \\
&+ (1-p_0)[\int_0^s \beta^B e^{-\beta^B\nu}e^{-\rho_1\nu}\,\mathrm{d}\nu + e^{-\rho_1 s}e^{-\beta^B s}]m.
\end{aligned}$$

Hence,

$$\begin{aligned}
\frac{\partial u^{\boldsymbol{r}_s,\boldsymbol{x}_M}(p_0)}{\partial s} &= p_0 e^{-(\rho_1+\gamma)s}(\rho_1+\gamma)(h-m) - (1-p_0)e^{-(\rho_1+\beta^B)s}\rho_1 m \\
&= (1-p_0)e^{-(\rho_1+\beta^B)s}[\frac{p_0}{1-p_0}e^{-(\gamma-\beta^B)s}(\rho_1+\gamma)(h-m) - \rho_1 m] \\
&= (1-p_0)e^{-(\rho_1+\beta^B)s}[\frac{p_s}{1-p_s}(\rho_1+\gamma)(h-m) - \rho_1 m].
\end{aligned}$$

Since $p_s \in [p_0, p_{fast}]$, and $\frac{p_{fast}}{1-p_{fast}} = \frac{\rho_1 m}{(\rho_1+\gamma)(h-m)}$, we have $\frac{\partial u^{\boldsymbol{r}_s,\boldsymbol{x}_M}(p_0)}{\partial s} < 0$ for $s \in [0, \hat{s})$. Hence, it is optimal for the evaluator to dismiss the agent immediately when $p_0 < p_{fast}$.

Hence, the best response of the evaluator to the strategy $\boldsymbol{x}_M$ is a cutoff strategy with a cutoff at p_{fast}.

Now, let us solve the best response of the type B agent given the evaluator's mixed cutoff strategy with a constant hazard rate $z^* = \frac{w}{c} - \rho_0 - \beta^B$ at the cutoff p_{fast}.

Let this strategy be $\boldsymbol{r}_c$. Clearly, since the evaluator would overthrow the agent when $p_t < p_{fast}$, thus the agent would never censor any bad news when $p_t < p_{fast}$. When $p_t = p_{fast}$, the evaluator faces a stationary problem since the evaluator would dismiss the agent at a constant rate. Hence,

$$V_B^{\boldsymbol{r}_c}(\bar{\varnothing}_t^B) = \max\left\{-\rho_0 c + \mathbb{E}[\int_0^{\lambda\wedge\tau^B} \rho_0 e^{-\rho_0\nu} w \,\mathrm{d}\nu + e^{-\rho_0\tau^B} V_B^{\lambda}(\varnothing_{\tau^B}^b)\,\mathbb{1}_{\{\tau^B<\lambda\}}], 0\right\},$$

where λ is the arrival time of dismissal induced by $z^* = \frac{w}{c} - \rho_0 - \beta^B$.

Since $\mathbb{E}[\int_0^{\lambda\wedge\tau^B} \rho_0 e^{-\rho_0\nu} w \,\mathrm{d}\nu] = \rho_0 c$, clearly the type B agent would be indifferent between censoring or nor. This also means that the continuation value of the type B agent is $\rho_0 c$ when no bad news arrives at $p_t = p_{fast}$.

When $p_t > p_{fast}$, the evaluator believes that the evaluator would censor all bad news until the public belief reaches p_{fast}. Hence, the type B agent can stay in power for some positive time $s = \frac{\ln[\frac{p_t}{1-p_t}\frac{1-p_{fast}}{p_{fast}}]}{\gamma} > 0$. Censoring all bad news until then gives her a payoff

$$\begin{aligned}
v_B^{\boldsymbol{r}_c,\boldsymbol{x}_M}(\bar{\varnothing}_t^B) &= -\rho_0 c + \int_0^s \rho_0 e^{-\rho_0\nu}(w - \beta^B c)\,\mathrm{d}\nu + e^{-\rho_0 s}\rho_0 c \\
&= [w - (\rho_0 + \beta^B)c](1 - e^{-\rho_0 s}) > 0.
\end{aligned}$$

Hence it is optimal to censor all bad news when $p_t > p_{fast}$. Thus, the strategy $\boldsymbol{x}_M$ is the best response to the evaluator's strategy, and the PBE is verified.

(2) Now we show the above PBE is the unique cutoff equilibrium. First, we show that for any admissible strategy of the agent, the evaluator's best response, if it exists, is to experiment when $p_t > p_{fast}$ and to dismiss the agent when $p_t < p_{slow}$. This helps

us restrict our attention to the evaluator's cutoff strategy $\hat{p} \leq p_{fast}$. Second, we show some necessary conditions for the agent's equilibrium strategy in a cutoff equilibrium. At last, we combine the above results to show no other cutoff equilibrium exists.

(2a) As in the proof of Proposition 2, the maximal information (no censorship) and the minimal information (full censorship) provide two bounds for the value function of the evaluator for any admissible strategy of the agent.

From the Proposition 1 and Proposition 2, we know that in both extreme cases, the evaluator would experiment when $p_t > p_{fast}$, and would dismiss the agent when $p_t < p_{slow}$. Hence, for any admissible strategy, the evaluator would do the same.

(2b) Fix a cutoff equilibrium where the evaluator uses a cutoff strategy with the cutoff belief $\hat{p}$, then we show that the equilibrium strategy of the type B agent is to censor all bad news whenever $p_t > \hat{p}$, and when $p_t = \hat{p}$, her censoring probability must satisfy $x_t^B \geq x^{B*} := 1 - \frac{\gamma}{\beta}$.

Suppose the equilibrium strategy of the type B agent is $x_t^B < 1$ at time t when $p_t > \hat{p}$.

First, note that the public belief in the absence of news will eventually drift down to $\hat{p}$ in finite time $s = \inf\{\nu > t : p_\nu = \hat{p}\} > 0$, otherwise the type B ruler would censor all bad news since the cost c is low (i.e. $c < \underline{c}$), which is a contradiction. Hence, there must be some $t' \in (t, s)$ such that $\dot{p}_{t'} < 0$; otherwise the pubic belief cannot drift down to $\hat{p}$. Thus, $x_{t'}^B > 0$, and $V_B^{\boldsymbol{r}}(\bar{\varnothing}_{t'}^B) \geq 0$. The exactly same argument in Proposition 2 shows that an alternative strategy $\boldsymbol{x}^{B'}$ such that censor all bad news for time $\nu \in [t, t']$ and resume the original strategy gives $v^{\boldsymbol{r}, \boldsymbol{x}^{B'}}(\bar{\varnothing}_t^B) > 0$. Hence, it is optimal for the type B agent to censor bad news at time t, i.e. $x_t^B = 1$. Contradiction.

Suppose the equilibrium strategy of the type B agent is $x_t^B < x^{B*}$ at time t when $p_t = \hat{p}$.

Clearly, when $x_t^B < 1 - \frac{\gamma}{\beta}$ and $p_t = \hat{p}$ in equilibrium, then $\dot{p}_t > 0$. Hence, for continuation games, $\dot{p}_\nu > 0$ for ν in some half-neighborhood of t, say $[t, t')$, since x_t^B is continuous from the right. It is clear that for any $\nu \in (t, t')$, $p_\nu > p^g$. Hence, by the above result, $x_\nu^B = 1$ for $\nu \in (t, t')$, which implies that $\dot{p}_\nu < 0$ for $\nu \in (t, t')$. Contradiction.

(2c) Suppose there is another cutoff equilibrium, in which the evaluator uses a cutoff strategy with a cutoff belief $\hat{p}$.

According to (2a), $\hat{p} \leq p_{fast}$. (2b) shows that in the equilibrium, the type B agent would censor all bad news when $p_t > \hat{p}$, and her censoring probability when $p_t = \hat{p}$ cannot be less than x^{B*}. Moreover, her censoring probability when $p_t = \hat{p}$ cannot be more than x^{B*} either. Suppose not, i.e. suppose $p_t = \hat{p}$ at some time t, and the type B agent's censoring probability is larger than x^{B*}. Then, the public belief would be drifting down below $\hat{p}$ after t. Since the evaluator uses a cutoff strategy, he must remove the agent at time t. However, the type B agent would not censor bad news at all at time t. This is a contradiction. Hence, in this cutoff equilibrium, the censoring probability when $p_t = \hat{p}$ must be equal to $x^{B*} \in (0, 1)$, and the public belief would stay at $\hat{p}$ thereafter. This means the evaluator's strategy at this belief must make the type B agent indifferent between censoring and not censoring bad news. This implies that the evaluator would use a constant hazard rate z^* to remove the agent at this belief. When $\hat{p} = p_{fast}$, this is the cutoff equilibrium we verified. When $\hat{p} < p_{fast}$, we can show that the best response of the evaluator is actually a cutoff strategy with the cutoff belief p_{fast}, since more censorship weakly increases the cutoff belief in the evaluator's cutoff strategy. This contradiction completes the proof.

□

Proof of Proposition 5. The first result is obvious.

When $p_0 \geq p_{fast}$, the type B agent's expected payoff in the equilibrium is

$$(w - \beta^B c)(1 - e^{-\rho_0 \hat{s}}) + e^{-\rho_0 \hat{s}} \rho_0 c.$$

Her payoff in the NCB is

$$\frac{\rho_0 w}{\rho_0 + \beta^B}.$$

Hence, she would be worse off in the equilibrium with censorship if and only if

$$(w - \beta^B c)(1 - e^{-\rho_0 \hat{s}}) + e^{-\rho_0 \hat{s}} \rho_0 c < \frac{\rho_0 w}{\rho_0 + \beta^B}.$$

Since we have

$$w - \beta^B c > \frac{\rho_0 w}{\rho_0 + \beta^B} > \rho_0 c,$$

thus she would be worse off in the equilibrium with censorship if and only if

$$e^{-\rho_0 \hat{s}} > \frac{\beta^B}{\rho_0 + \beta^B}.$$

Denote $\hat{s}^* := \frac{1}{\rho_0} \ln[\frac{\rho_0 + \beta^B}{\beta^B}]$, we get the result.

□

Proof of Corollary 3. Clearly, since $\hat{s}$ is not a function of the censoring cost c, whether the type B agent is better off in the equilibrium than she is in the NCB does not depend on the cost.

when $p_0 \geq p_{fast}$, the type B agent's expected payoff in the equilibrium is

$$(w - \beta^B c)(1 - e^{-\rho_0 \hat{s}}) + e^{-\rho_0 \hat{s}} \rho_0 c.$$

Its first derivative with respect to c is

$$-\beta^B + (\rho_0 + \beta^B) e^{-\rho_0 \hat{s}}.$$

It is clear that the type B agent's expected payoff in the equilibrium is decreasing/increasing in the censoring cost c if and only if

$$e^{-\rho_0 \hat{s}} \lesseqgtr \frac{\beta^B}{\rho_0 + \beta^B}.$$

From the proof of Proposition 5, clearly, the conditions coincide.

□

Proof of Proposition 6. (1) First, given the strategies of all players, the pubic belief will drift and jump differently in different phases. When $p_t > p^B$, no bad news is expected in equilibrium, and we assume the public belief will jump to 0 after off-path bad news. In addition, the drifting process in the absence of news follows

$$\mathrm{d}p_t = -p_t(1 - p_t)\gamma \, \mathrm{d}t.$$

When $p_t \in (p^G, p^B)$, the public belief will jump to $J(p_t, 1, 0) = 0$ after a piece of bad news, and the drifting process in the absence of news follows

$$\mathrm{d}p_t = -p_t(1 - p_t)(\gamma - \beta^B) \, \mathrm{d}t.$$

When $p_t < p^G$, the public belief will jump to $j(p_t) > 0$ after a piece of bad news, and the drifting process in the absence of news follows

$$\mathrm{d}p_t = -p_t(1 - p_t)(\gamma + \beta^G - \beta^B) \, \mathrm{d}t.$$

In addition, we will see later that $j(p^G) \le p_{fast}$ is equivalent to $\beta^G \le \bar{\beta}$, hence $j(p_t) \le j(p^G) \le p_{fast}$ when $p_t < p^G$.

Hence, it means that the public belief p_t will jump below p_{fast} after a piece of bad news, no matter how high $p_t < 1$ is.

(2) Now, suppose that the evaluator uses a cutoff strategy with the cutoff p_{fast}, let us solve the best response of each type of the agent.

Since the type B agent faces the exactly same problem as in the conclusive bad news case. Her optimal strategy is the cutoff strategy with the cutoff belief p^B.

The logic in Proposition 2 also applies to the type G agent. She will eventually stop censoring when the public belief approaches p_{fast}. Let τ^G and χ be the arrival times of a piece of bad news and a piece of good news (i.e. success) from the type G agent, respectively. Then her payoff after stopping censoring, given the public belief $p > p_{fast}$, is

$$\mathbb{E}[\int_0^T \rho_0 e^{-\rho_0 \nu} w d\nu] = \mathbb{E}[w(1 - e^{-\rho_0 T})],$$

where

$$T = \begin{cases} \infty, & \chi < \tau^G \wedge s, \\ \tau^G \wedge s, & \chi > \tau^G \wedge s, \end{cases}$$

and s is the time when the public belief drifts from p to p_{fast}. We can show that

$$\mathbb{E}[w(1-e^{-\rho_0 T})] = \frac{\gamma+\rho_0}{\beta^G+\gamma+\rho_0}(1-e^{-(\beta^G+\gamma+\rho_0)s})w.$$

Its partial derivative with respect to s is

$$(\gamma+\rho_0)e^{-(\beta^G+\gamma+\rho_0)s}w > 0.$$

Hence, it is strictly increasing in s.

The same argument in Proposition 2 also implies that, at public belief p^G, the type G agent must be indifferent between censoring or not if a piece of bad news arrives,

$$\frac{\gamma+\rho_0}{\beta^G+\gamma+\rho_0}(1-e^{-(\beta^G+\gamma+\rho_0)s_G})w = \rho_0 c,$$

where $s_G = \frac{\ln[\frac{p^G}{1-p^G}\frac{1-p_{fast}}{p_{fast}}]}{\gamma+\beta^G-\beta^B}$. This indifference condition pins down p^G. The type G agent has a strict incentive to censor all bad news when $p_t > p^G$, and has a strict incentive not to censor any news when $p_t < p^G$.

Let s_B be such that

$$\mathbb{E}[-\rho_0 c + \int_0^{s_B \wedge \tau^B} \rho_0 e^{-\rho_0 \nu} w \, \mathrm{d}\nu] = 0.$$

Hence, s_B is the time when the public belief drifts from $p^{B\dagger}$ to p_{fast}. Thus

$$\frac{\rho_0}{\beta^B+\rho_0}(1-e^{-(\beta^B+\rho_0)s_B})w = \rho_0 c.$$

Since $c < \underline{c}$, we have

$$\rho_0 c < \frac{\rho_0}{\beta^B+\rho_0}w < \frac{\gamma+\rho_0}{\beta^G+\gamma+\rho_0}w.$$

It is easy to see that $\frac{\gamma+\rho_0}{\beta^G+\gamma+\rho_0}(1-e^{-(\beta^G+\gamma+\rho_0)s})w$ is increasing in $\gamma > 0$, and decreasing in $\beta^G < \beta^B$, thus we have

$$\begin{aligned}\frac{\rho_0}{\beta^B+\rho_0}(1-e^{-(\beta^B+\rho_0)s_B})w &= \frac{\gamma+\rho_0}{\beta^G+\gamma+\rho_0}(1-e^{-(\beta^G+\gamma+\rho_0)s_G})w \\ &> \frac{\rho_0}{\beta^B+\rho_0}(1-e^{-(\beta^B+\rho_0)s_G})w.\end{aligned}$$

Hence, $s_B > s_G$, which implies $p^{B\dagger} > p^G$.

Now we show that there exists a $\bar{\beta} < \beta^B$ such that $j(p^G) \leq p_{fast}$ if and only if $\beta^G < \bar{\beta}$.

Note that

$$j^{-1}(p_{fast}) = \frac{p_{fast}\beta^B}{p_{fast}\beta^B + (1-p_{fast})\beta^G} = \frac{p_{fast}\frac{\beta^B}{\beta^G}}{p_{fast}\frac{\beta^B}{\beta^G} + (1-p_{fast})},$$

and

$$p^G = \frac{p_{fast}e^{(\gamma+\beta^G-\beta_B)s_G}}{1-p_{fast}+p_{fast}e^{(\gamma+\beta^G-\beta_B)s_G}}.$$

Hence, $j^{-1}(p_{fast}) \geq p^G$ is equivalent to $\frac{\beta^B}{\beta^G} \geq e^{(\gamma+\beta^G-\beta_B)s_G}$. Since

$$\rho_0 c = \frac{\gamma+\rho_0}{\beta^G+\gamma+\rho_0}(1-e^{-(\beta^G+\gamma+\rho_0)s_G})w,$$

$j^{-1}(p^g) \geq p_G^*$ is also equivalent to

$$\frac{\beta^B}{\beta^G} \geq (\frac{\bar{c}}{\bar{c}-c})^{1-\frac{\beta^B+\rho_0}{\gamma+\beta^G+\rho_0}},$$

where

$$\bar{c} = \frac{\gamma+\rho_0}{\beta^G+\gamma+\rho_0}\frac{w}{\rho_0}.$$

Note that

$$\frac{\mathrm{d}\beta^G(\frac{\bar{c}}{\bar{c}-c})^{1-\frac{\beta^B+\rho_0}{\gamma+\beta^G+\rho_0}}}{\mathrm{d}\beta^G} = (\frac{\bar{c}}{\bar{c}-c})^{1-\frac{\beta^B+\rho_0}{\gamma+\beta^G+\rho_0}}\{1+\frac{\beta^G}{\gamma+\beta^G+\rho_0}[\frac{\rho_0 c(\gamma+\beta^G-\beta^B)}{w(\gamma+\rho_0)-\rho_0 c(\gamma+\beta^G+\rho_0)} + \ln(\frac{\bar{c}}{\bar{c}-c})\frac{\beta^G+\rho_0}{\gamma+\beta^G+\rho_0}]\} > 0.$$

In addition, $\beta^G(\frac{\bar{c}}{\bar{c}-c})^{1-\frac{\beta^B+\rho_0}{\gamma+\beta^G+\rho_0}}$ goes to 0 when β^G goes to 0, and it is lager than β^B when $\beta^G = \beta^B$. Hence, there exists a $\bar{\beta} \in (0, \beta^B)$, such that when $\beta^G \in (0, \bar{\beta}]$, the condition $j^{-1}(p_{fast}) \geq p^G$ is satisfied.

(3) At last, given the strategies of both types of agent, the evaluator faces the classic bandit problem when $p_t < p^G$, and his best response is a cutoff strategy with the cutoff p_{fast}. In addition, when $p_t > p^G$, it is easy to verify that the evaluator has a strict incentive to not dismiss the agent. This completes the proof. □

Proof of Proposition 7. (1) We first establish that there exists a $\underline{\beta} \in (0, \bar{\beta})$, such that $j(p^{B\dagger}) \leq p_{fast}$ if and only if $\beta^G \in (0, \underline{\beta})$.

Note that

$$p^{B\dagger} = \frac{\frac{p^G}{1-p^G}e^{(\gamma-\beta^B)(s_B-s_G)}}{1+\frac{p^G}{1-p^G}e^{(\gamma-\beta^B)(s_B-s_G)}} = \frac{\frac{p_{fast}}{1-p_{fast}}e^{(\gamma-\beta^B)s_B}e^{\beta^G s_G}}{1+\frac{p_{fast}}{1-p_{fast}}e^{(\gamma-\beta^B)s_B}e^{\beta^G s_G}} = \frac{p_{fast}e^{(\gamma-\beta^B)s_B}e^{\beta^G s_G}}{1-p_{fast}+p_{fast}e^{(\gamma-\beta^B)s_B}e^{\beta^G s_G}}.$$

Hence, $j^{-1}(p_{fast}) \geq p^{B\dagger}$ is equivalent to $\frac{\beta^B}{\beta^G} \geq e^{(\gamma-\beta^B)s_B}e^{\beta^G s_G}$. Since

$$\rho_0 c = \frac{\gamma+\rho_0}{\beta^G+\gamma+\rho_0}(1-e^{-(\beta^G+\gamma+\rho_0)s_G})w,$$

we have

$$\beta^G s_G = \beta^G \frac{\ln \frac{w(\gamma+\rho_0)}{w(\gamma+\rho_0)-\rho_0 c(\beta^G+\gamma+\rho_0)}}{\beta^G+\gamma+\rho_0}.$$

Clearly, $\beta^G s_G$ is increasing in β^G.

Consider

$$\frac{\beta^B}{\beta^G} \geq e^{(\gamma-\beta^B)s_B} e^{\beta^G s_G}.$$

The left hand side (LHS) is decreasing in β^G, and the right hand side (RHS) is increasing in β^G. In addition, $\beta^G e^{(\gamma-\beta^B)s_B} e^{\beta^G s_G}$ goes to 0 when β^G goes to 0, and it is larger than β^B when $\beta^G = \beta^B$.

Thus, $j^{-1}(p_{fast}) \geq p^{B\dagger}$ if and only if β^G is smaller than some threshold $\underline{\beta} > 0$, and $\underline{\beta} < \bar{\beta}$ since $p^{B\dagger} > p^G$.

(2) We show that the evaluator is better off with censorship when $p_0 \in (p^G, p^{B\dagger}]$.

For $p \in [p^G, p^{B\dagger}]$, the HJB equation in the NCB is

$$(\rho_1+p\gamma)U^{\mathbf{0}}(p)+(p\beta^G+(1-p)\beta^B)[U^{\mathbf{0}}(p)-U^{\mathbf{0}}(j(p))]+(\gamma+\beta^G-\beta^B)p(1-p)U^{\mathbf{0}\prime}(p) = p\gamma k(\gamma+\rho_1).$$

In equilibrium, it is

$$(\rho_1+p\gamma)U^{\tilde{\mathbf{x}}}(p)+(1-p)\beta^B[U^{\tilde{\mathbf{x}}}(p)-m]+(\gamma-\beta^B)p(1-p)U^{\tilde{\mathbf{x}}\prime}(p) = p\gamma k(\gamma+\rho_1).$$

For $p \leq p^G$, clearly $U^{\tilde{\mathbf{x}}}(p) = U^{\mathbf{0}}(p)$. Hence, at $p = p^G$, comparing the two HJB equations, we have

$$U^{\mathbf{0}\prime}(p^G) < U^{\tilde{\mathbf{x}}\prime}(p^G),$$

where $U^{\tilde{\mathbf{x}}\prime}(p^G)$ is the right derivative.

Since the value functions are continuously differentiable, the above relation is true for some neighborhood $p \in [p^G, \check{p})$, hence in that neighborhood $U^{\tilde{\mathbf{x}}}(p) - U^{\mathbf{0}}(p)$ is increasing in p. Thus, $U^{\tilde{\mathbf{x}}}(p) > U^{\mathbf{0}}(p)$ for $p \in (p^G, \check{p})$.

To prove by contradiction, suppose there is some $\acute{p} \in (p^G, p^{B\dagger}]$ such that $U^{\tilde{\mathbf{x}}}(\acute{p}) \leq U^{\mathbf{0}}(\acute{p})$. Let $\acute{p} = \inf\{p \in (p^G, p^{B\dagger}] : U^{\tilde{\mathbf{x}}}(\acute{p}) = U^{\mathbf{0}}(\acute{p})\}$. First, note that it must be that $U^{\mathbf{0}\prime}(\acute{p}) \geq U^{\tilde{\mathbf{x}}\prime}(\acute{p})$. Otherwise, if $U^{\mathbf{0}\prime}(\acute{p}) < U^{\tilde{\mathbf{x}}\prime}(\acute{p})$, then it is true in a small neighborhood, and $U^{\tilde{\mathbf{x}}}(p) - U^{\mathbf{0}}(p)$ is increasing in p in that neighborhood, then it must be that $U^{\tilde{\mathbf{x}}}(\acute{p}) > U^{\mathbf{0}}(\acute{p})$. Then

$$\begin{aligned}
&(\acute{p}\beta^G+(1-\acute{p})\beta^B)[U^{\mathbf{0}}(\acute{p})-U^{\mathbf{0}}(j(\acute{p}))]+(\gamma+\beta^G-\beta^B)\acute{p}(1-\acute{p})U^{\mathbf{0}\prime}(\acute{p}) \\
=&(1-\acute{p})\beta^B[U^{\tilde{\mathbf{x}}}(\acute{p})-m]+(\gamma-\beta^B)\acute{p}(1-\acute{p})U^{\tilde{\mathbf{x}}\prime}(\acute{p}) \\
\leq&(1-\acute{p})\beta^B[U^{\tilde{\mathbf{x}}}(\acute{p})-m]+(\gamma-\beta^B)\acute{p}(1-\acute{p})U^{\mathbf{0}\prime}(\acute{p}).
\end{aligned}$$

Hence,

$$(\acute{p}\beta^G + (1-\acute{p})\beta^B)[U^{\mathbf{0}}(\acute{p}) - U^{\mathbf{0}}(j(\acute{p}))] + \beta^G \acute{p}(1-\acute{p})U^{\mathbf{0}\prime}(\acute{p}) - (1-\acute{p})\beta^B[U^{\mathbf{0}}(\acute{p}) - m] \leq 0.$$

Note that if $j(\acute{p}) \leq p_{fast}$, then $U^{\mathbf{0}}(j(\acute{p})) = m$, a contradiction. Hence, it must be $j(\acute{p}) > p_{fast}$. Since $U^{\mathbf{0}}(p)$ is strictly convex in $p \in [p_{fast}, 1]$, we have

$$U^{\mathbf{0}\prime}(\acute{p}) > \frac{U^{\mathbf{0}}(\acute{p}) - U^{\mathbf{0}}(j(\acute{p}))}{\acute{p} - j(\acute{p})},$$

and

$$U^{\mathbf{0}}(\acute{p}) - m = U^{\mathbf{0}}(\acute{p}) - U^{\mathbf{0}}(p_{fast}) < [U^{\mathbf{0}}(\acute{p}) - U^{\mathbf{0}}(j(\acute{p}))]\frac{\acute{p} - p_{fast}}{\acute{p} - j(\acute{p})}.$$

Hence, we have

$$\begin{aligned}
&(\acute{p}\beta^G + (1-\acute{p})\beta^B)[U^{\mathbf{0}}(\acute{p}) - U^{\mathbf{0}}(j(\acute{p}))] + \beta^G \acute{p}(1-\acute{p})U^{\mathbf{0}\prime}(\acute{p}) - (1-\acute{p})\beta^B[U^{\mathbf{0}}(\acute{p}) - m] \\
>&(\acute{p}\beta^G + (1-\acute{p})\beta^B)[U^{\mathbf{0}}(\acute{p}) - U^{\mathbf{0}}(j(\acute{p}))] + \beta^G \acute{p}(1-\acute{p})\frac{U(\acute{p}) - U(j(\acute{p}))}{\acute{p} - j(\acute{p})} \\
&- (1-\acute{p})\beta^B[U^{\mathbf{0}}(\acute{p}) - U^{\mathbf{0}}(j(\acute{p}))]\frac{\acute{p} - p_{fast}}{\acute{p} - j(\acute{p})} \\
=&(\acute{p}\beta^G + (1-\acute{p})\beta^B + \beta^G \acute{p}(1-\acute{p})\frac{1}{\acute{p} - j(\acute{p})} - (1-\acute{p})\beta^B\frac{\acute{p} - p_{fast}}{\acute{p} - j(\acute{p})})[U^{\mathbf{0}}(\acute{p}) - U^{\mathbf{0}}(j(\acute{p}))] \\
>&(\acute{p}\beta^G + (1-\acute{p})\beta^B + \beta^G \acute{p}(1-\acute{p})\frac{1}{\acute{p} - j(\acute{p})} - (1-\acute{p})\beta^B\frac{\acute{p}}{\acute{p} - j(\acute{p})})[U^{\mathbf{0}}(\acute{p}) - U^{\mathbf{0}}(j(\acute{p}))] \\
=&(\acute{p}\beta^G + (1-\acute{p})\beta^B + \frac{\acute{p}(1-\acute{p})(\beta^G - \beta^B)}{\acute{p} - j(\acute{p})})[U^{\mathbf{0}}(\acute{p}) - U^{\mathbf{0}}(j(\acute{p}))].
\end{aligned}$$

Since

$$\acute{p} - j(\acute{p}) = \frac{(\beta^B - \beta^G)\acute{p}(1-\acute{p})}{\acute{p}\beta^G + (1-\acute{p})\beta^B},$$

the above equation is strictly larger than 0, a contradiction.

(3) Now we show the type B agent is better off with censorship when $\beta^G \in (0, \underline{\beta}]$ and $p_0 \in (p^G, p^{B\dagger}]$.

In the equilibrium and the NCB, the type B agent does not censor bad news when $p_0 \leq p^{B\dagger}$. In addition, in both cases, she would be dismissed when either a piece of bad news arrives or when the public belief drift down to p_{fast}. Since the public belief drifts down slower in the equilibrium than it does in the NCB, the type B agent is better off with censorship.

(4) At last, we show the type G agent is also better off with censorship when $\beta^G \in (0, \underline{\beta}]$ and $p_0 \in (p^G, p^{B\dagger}]$.

In the equilibrium, her payoff when $p_0 \in (p^G, p^{B\dagger}]$ is

$$\begin{aligned}&\int_0^{s^1} \gamma e^{-\gamma\nu}[w - c\int_0^{\nu} \rho_0 e^{-\rho_0 t}\beta^G \,\mathrm{d}t] + e^{-\gamma s^1}[(w - c\beta^G)(1 - e^{-\rho_0 s^1}) + e^{-\rho_0 s^1}\rho_0 c] \\ =&\frac{\gamma w + \rho_0(w - c\beta^G)}{\gamma + \rho_0}(1 - e^{-(\rho_0+\gamma)s^1}) + e^{-(\rho_0+\gamma)s^1}\rho_0 c,\end{aligned}$$

where $s^1 = \frac{\ln[\frac{p_0}{1-p_0}\frac{1-p^G}{p^G}]}{\gamma-\beta^B}$ is the time that belief drifts down from p_0 to p^G according to rate $\gamma - \beta^B$.

In the NCB, since $j(p_0) \leq p_{fast}$, her payoff is

$$\begin{aligned}&\int_0^{s^0} \beta^G e^{-\beta^G\nu} e^{-\gamma\nu} w(1 - e^{-\rho_0\nu})d\nu + \int_0^{s^0} \gamma e^{-\beta^G\nu} e^{-\gamma\nu} w d\nu + e^{-\beta^G s^0} e^{-\gamma s^0}[w(1 - e^{-\rho_0 s^0}) + e^{-\rho_0 s^0}\rho_0 c] \\ =&\frac{\gamma + \rho_0}{\beta^G + \gamma + \rho_0}(1 - e^{-(\beta^G+\gamma+\rho_0)s^0})w + e^{-(\beta^G+\gamma+\rho_0)s^0}\rho_0 c,\end{aligned}$$

where $s^0 = \frac{\ln[\frac{p_0}{1-p_0}\frac{1-p^G}{p^G}]}{\gamma+\beta^G-\beta^B}$ is the time that belief drifts down from p_0 to p^G according to rate $\gamma + \beta^G - \beta^B$.

First, since $c < \underline{c}$, we have

$$\frac{\gamma w + \rho_0(w - c\beta^G)}{\gamma + \rho_0} > \frac{\gamma + \rho_0}{\beta^G + \gamma + \rho_0} w.$$

Second, note that

$$e^{-(\rho_0+\gamma)s^1} < e^{-(\beta^G+\gamma+\rho_0)s^0},$$

because

$$(\rho_0+\gamma)s^1 = \ln[\frac{p_0}{1-p_0}\frac{1-p^G}{p^G}]\frac{\rho_0+\gamma}{\gamma-\beta^B} > \ln[\frac{p_0}{1-p_0}\frac{1-p^G}{p^G}]\frac{\beta^G+\gamma+\rho_0}{\gamma+\beta^G-\beta^B} = (\beta^G+\gamma+\rho_0)s^0.$$

Hence, the type G agent is better off with censorship.

□

Proof of Proposition 8. (1) (Low cost) Given the strategies of both agents, there is no bad news in equilibrium when $p_t > p_{fast}$. We assume the public belief jumps down to 0 after off-path bad news. In addition, the belief drifting process is the same as in Proposition 4 when $p_t \geq p_{fast}$.

Given the strategy of the type G agent, the evaluator and the type B agent face the same problem, hence have the same best response as in Proposition 4 when $p_t \geq p_{fast}$.

Given the strategy of the evaluator, when $p_t = p_{fast}$, the type G agent's continuation value after censoring one piece of bad news is

$$\mathbb{E}[\int_0^T \rho_0 e^{-\rho_0 \nu} w \, \mathrm{d}\nu - \int_0^{T\wedge\chi} \rho_0 e^{-\rho_0 \nu} c X_\nu^G \, \mathrm{d}N_\nu^G],$$

where

$$T = \begin{cases} \infty, & \chi < \lambda, \\ \lambda, & \chi > \lambda, \end{cases}$$

and λ is the arrival time of dismissal induced by $z^* = \frac{w}{c} - \rho_0 - \beta^B$, and χ is the arrival time of success.

Hence, her continuation value is

$$\begin{aligned} &\mathbb{E}[w(1 - e^{-\rho_0 T}) - \beta^G c(1 - e^{-\rho_0 (T\wedge\chi)})] \\ =& w\frac{\gamma + \rho_0}{z^* + \gamma + \rho_0} - \beta^G c \frac{\rho_0}{z^* + \gamma + \rho_0} \\ =& \frac{w(\gamma + \rho_0) - \beta^G \rho_0 c}{z^* + \gamma + \rho_0}. \end{aligned}$$

Since $z^* = \frac{w}{c} - \rho_0 - \beta^B$, we have

$$\frac{w(\gamma + \rho_0) - \beta^G \rho_0 c}{z^* + \gamma + \rho_0} > \rho_0 \frac{w - \beta^G c}{z^* + \rho_0} = \rho_0 c \frac{w - \beta^G c}{w - \beta^B c} > \rho_0 c.$$

Hence, the type G agent has a strict incentive to censor bad news when $p_t = p_{fast}$. In addition, it is easy to verify that she also has a strict incentive to censor bad news when $p_t > p_{fast}$.

At last, when $p_t < p_{fast}$, we can show that the evaluator has a strict incentive to dismiss the agent, given the strategies of both types of agent. Thus, neither agent has an incentive to censor bad news when $p_t < p_{fast}$.

(2) (Intermediate cost) Given the strategies of both agents, a piece of bad news will make the public belief jump down to $J(p_t, 1, 0) = 0$ when $p_t > p_{slow}$. In the absence of news, the public belief will drift up according to

$$\mathrm{d}p_t = -p_t(1 - p_t)(\gamma - \beta^B) \, \mathrm{d}t.$$

The type B agent faces a problem as in Proposition 1, since $c > \underline{c}$, so she never censors any bad news.

The evaluator also faces a problem as in Proposition 1, thus he uses a cutoff strategy with the cutoff belief p_{slow}.

When $p_t > p_{slow}$, if the type G agent censors all bad news before the success, her payoff would be

$$\int_0^\infty \gamma e^{-\gamma\nu}[w - c\int_0^\nu \rho_0 e^{-\rho_0 t}\beta^G \,\mathrm{d}t]\,\mathrm{d}\nu = w - c\beta^G + c\beta^G\frac{\gamma}{\rho_0+\gamma} > \rho_0 c,$$

since $c < \bar{c}$. Hence, the type G agent has a strict incentive to censor bad news.

At last, when $p_t < p_{slow}$, we can show that the evaluator has a strict incentive to dismiss the agent, given the strategies of both types of agent. Thus, neither agent has an incentive to censor bad news when $p_t < p_{slow}$. □

Proof of Proposition 9. (1) First, we summarize the property of p^* – the strategy of the evaluator in the NCB from the bandit literature. The optimal policy of the evaluator is a cutoff strategy with the cutoff p^*. Her value function is continuously differentiable everywhere, with a possible exception at the cutoff p^*. In addition, when $p_0 > p^*$, her value function is strictly increasing and strictly convex.

When $\gamma + \beta^G \geq \beta^B$, $p^* = p_{fast}$, and the evaluator's value function is continuously differentiable at p^*.

When $\gamma + \beta^G < \beta^B$, $p^* \in (p_{slow}, p_{fast})$, and the evaluator's value function has a kink at p^*.

(2) (Low cost) Assume $c < \underline{c}$ and $p_0 > p^*$.

In the NCB, the evaluator can always use the same strategy as in the equilibrium by ignoring bad news and dismissing the agent if no success arrives for some time. However, this strategy is strictly dominated by her optimal strategy in the NCB, which implies she is strictly worse off in the equilibrium when $p_0 > p^*$.

(3) (Intermediate cost) Assume $c \in (\underline{c}, \bar{c})$ and $p_0 > p_{slow}$. Let $U^{\mathbf{0}}(p_0)$ and $U^{\tilde{\mathbf{x}}}(p_0)$ be the evaluator's value functions in the NCB and in the equilibrium, respectively.

Clearly, the evaluator is better off in equilibrium when $p_0 \in (p_{slow}, p^*]$ and $U^{\mathbf{0}}(p^*) < U^{\tilde{\mathbf{x}}}(p^*)$.

Consider $p_0 \in (p^*, 1)$. To prove by contradiction, suppose there is some $\acute{p} \in (p^*, 1)$ such that $U^{\tilde{\mathbf{x}}}(\acute{p}) \leq U^{\mathbf{0}}(\acute{p})$. Let $\acute{p} = \inf\{p \in (p^*, 1) : U^{\tilde{\mathbf{x}}}(\acute{p}) = U^{\mathbf{0}}(\acute{p})\}$. First, note that it must be that $U^{\mathbf{0}\prime}(\acute{p}) \geq U^{\tilde{\mathbf{x}}\prime}(\acute{p})$. Otherwise, if $U^{\mathbf{0}\prime}(\acute{p}) < U^{\tilde{\mathbf{x}}\prime}(\acute{p})$, then it is true in a small neighborhood, and $U^{\tilde{\mathbf{x}}}(p) - U^{\mathbf{0}}(p)$ is increasing in p in that neighborhood, then it must be that $U^{\tilde{\mathbf{x}}}(\acute{p}) > U^{\mathbf{0}}(\acute{p})$, a contradiction.

Since $U^{\mathbf{0}}(p)$ is strictly convex for $p \in [p^*, 1]$, $U^{\mathbf{0}\prime}(p) > U^{\mathbf{0}\prime}(\acute{p}) \geq U^{\tilde{\mathbf{x}}\prime}(\acute{p})$ for $p > \acute{p}$. Also, from the Proof of Proposition 1, we know that the value function of the evaluator in the equilibrium is linear when $p_0 > p_{slow}$, hence $U^{\mathbf{0}\prime}(p) > U^{\tilde{\mathbf{x}}\prime}(\acute{p}) = U^{\tilde{\mathbf{x}}\prime}(p)$ for any $p \in (\acute{p}, 1)$. Hence, $U^{\mathbf{0}}(p) - U^{\tilde{\mathbf{x}}}(p)$ is continuous and strictly increasing in $p \in (\acute{p}, 1]$. Thus, $U^{\mathbf{0}}(1) - U^{\tilde{\mathbf{x}}}(1) > U^{\mathbf{0}}(\acute{p}) - U^{\tilde{\mathbf{x}}}(\acute{p}) = 0$, a contradiction.

□

References

Acharya, Viral V, Peter DeMarzo, and Ilan Kremer (2011), "Endogenous Information Flows and the Clustering of Announcements." *American Economic Review*, 101, 2955–2979.

Besley, Timothy and Andrea Prat (2006), "Handcuffs for the Grabbing Hand? Media Capture and Government Accountability." *The American Economic Review*, 96, 720–736.

Che, Yeon-Koo and Johannes Hörner (2018), "Recommender Systems as Mechanisms for Social Learning*." *The Quarterly Journal of Economics*, 133, 871–925.

Daughety, Andrew F. and Jennifer F. Reinganum (2018), "Evidence Suppression by Prosecutors: Violations of the Brady Rule." *SSRN Electronic Journal*.

Dye, Ronald A. (1985), "Disclosure of Nonproprietary Information." *Journal of Accounting Research*, 23, 123–145.

Dye, Ronald A. (2017), "Optimal disclosure decisions when there are penalties for nondisclosure." *The RAND Journal of Economics*, 48, 704–732.

Edmond, Chris (2013), "Information manipulation, coordination, and regime change." *Review of Economic Studies*, 80, 1422–1458.

Egorov, Georgy, Sergei Guriev, and Konstantin Sonin (2009), "Why resource-poor dictators allow freer media: A theory and evidence from panel data." *American Political Science Review*, 103, 645–668.

Ely, Jeffrey C. (2017), "Beeps." *American Economic Review*.

Eraslan, Hulya and Saltuk Ozerturk (2017), "Information Gatekeeping and Media Bias."

Gehlbach, Scott and Konstantin Sonin (2014), "Government control of the media." *Journal of Public Economics*, 118, 163–171.

Grossman, Sanford J. (1981), "The Informational Role of Warranties and Private Disclosure about Product Quality." *The Journal of Law & Economics*, 24, 461–483.

Grubb, Michael D. (2011), "Developing a reputation for reticence." *Journal of Economics & Management Strategy*, 20, 225–268.

Guriev, Sergei M. and Daniel Treisman (2018), "Informational Autocracy: Theory and Empirics of Modern Authoritarianism." Technical report.

Guttman, Ilan, Ilan Kremer, and Andrzej Skrzypacz (2014), "Not Only What but Also When: A Theory of Dynamic Voluntary Disclosure." *American Economic Review*, 104, 2400–2420.

Hauser, Daniel N. (2017), "Promotion, Censorship, and Reputation for Quality."

Jung, Woon-Oh and Young K. Kwon (1988), "Disclosure When the Market Is Unsure of Information Endowment of Managers." *Journal of Accounting Research*, 26, 146–153.

Kamenica, Emir and Matthew Gentzkow (2011), "Bayesian persuasion." *The American Economic Review*, 101, 2590–2615.

Kartik, Navin, Frances Xu Lee, and Wing Suen (2017), "A Theorem on Bayesian Updating and Applications to Communication Games."

Keller, Godfrey and Sven Rady (2010), "Strategic experimentation with Poisson bandits." *Theoretical Economics*, 5, 275–311.

Keller, Godfrey and Sven Rady (2015), "Breakdowns." *Theoretical Economics*, 10, 175–202.

Keller, Godfrey, Sven Rady, and Martin Cripps (2005), "Strategic Experimentation with Exponential Bandits." *Econometrica*, 73, 39–68.

Laraki, Rida, Eilon Solan, and Nicolas Vieille (2005), "Continuous-time games of timing." *Journal of Economic Theory*, 120, 206–238.

Li, Wei and Dennis Tao Yang (2005), "The great leap forward: Anatomy of a central planning disaster." *Journal of Political Economy*, 113, 840–877.

Lorentzen, Peter (2014), "China's Strategic Censorship." *American Journal of Political Science*, 58, 402–414.

Milgrom, Paul R. (1981), "Good News and Bad News: Representation Theorems and Applications." *The Bell Journal of Economics*, 12, 380–391.

Presman, E. L. (1991), "Poisson Version of the Two-Armed Bandit Problem with Discounting." *Theory of Probability & Its Applications*, 35, 307–317.

Redlicki, Jakub (2017), "What Drives Regimes to Manipulate Information: Criticism, Collective Action, and Coordination."

Renault, Jérôme, Eilon Solan, and Nicolas Vieille (2017), "Optimal dynamic information provision." *Games and Economic Behavior*, 104, 329–349.

Riedel, Frank and Jan-Henrik Steg (2017), "Subgame-perfect equilibria in stochastic timing games." *Journal of Mathematical Economics*, 72, 36–50.

Rodan, Garry (2004), *Transparency and Authoritarian Rule in Southeast Asia: Singapore and Malaysia*. Routledge.

Rosenbaum, Stephen Mark, Stephan Billinger, and Nils Stieglitz (2014), "Let's be honest: A review of experimental evidence of honesty and truth-telling." *Journal of Economic Psychology*, 45, 181–196.

Shadmehr, Mehdi and Dan Bernhardt (2015), "State censorship." *American Economic Journal: Microeconomics*, 7, 280–307.

Shin, Hyun Song (2003), "Disclosures and Asset Returns." *Econometrica*, 71, 105–133.

Sirowy, Larry and Alex Inkeles (1990), "The effects of democracy on economic growth and inequality: A review." *Studies in Comparative International Development*, 25, 126–157.

Smirnov, Aleksei and Egor Starkov (2018), "Bad News Turned Good: Reversal Under Censorship."