Motivational Ratings*

Johannes Hörner[†], Nicolas S. Lambert[‡]

September 3, 2016

Abstract

Rating systems not only provide information to users but also motivate the rated agent. This paper solves for the optimal (effort-maximizing) rating system within the standard career concerns framework. It is a mixture two-state rating system. That is, it is the sum of two Markov processes, with one that reflects the belief of the rater and the other the preferences of the rated agent. The rating, however, is not a Markov process. Our analysis shows how the rating combines information of different types and vintages. In particular, an increase in effort may affect some (but not all) future ratings adversely.

Keywords: Career Concerns; Mechanism Design; Ratings.

JEL codes: C72, C73

1 Introduction

Helping users make informed decisions is only one of the goals of ratings. Another is to motivate the rated firm or agent. These two goals are not necessarily aligned.

*We would like to thank Sümeyra Akin, Ian Ball, and Daria Khromenkova for excellent research assistance, as well as Drew Fudenberg, Marina Halac and David Kreps for their comments.

[†]Yale University, 30 Hillhouse Ave., New Haven, CT 06520, USA, johannes.horner@yale.edu.

[‡]Stanford Graduate School of Business, 655 Knight Way, Stanford, CA 94305, USA, nlambert@stanford.edu.

Excessive information depresses career concerns and distorts the agent's choices.[1] The purpose of this paper is to examine this trade-off. In particular, we ask the following: how should different sources of information be combined? At what rate, if any, should past observations be discounted? Finally, how do standard rating mechanisms compare?

We demonstrate that the optimal rating system always confounds the different signals yet never adds any irrelevant noise. To maximize incentives for effort, the rater combines the entire history of signals in a one-dimensional statistic, which neither is a simple function of the rater's current belief (about the agent's type) nor enables the market to back out this belief from the rating history. It is not simply a function of her latest rating and signal either.[2] Furthermore, the time series of ratings fails to satisfy the Markov property.[3]

However, the optimal rating system has a remarkably simple structure: it is a linear combination of two processes, namely, the rater's underlying belief and an *incentive state* that reflects both the agent's preferences and the determinants of the signal processes. That is, the optimal rating process admits a simple decomposition as a two-dimensional Markov mixture model.

The agent's preferences determine the *impulse response* of the incentive state via his impatience. That is, past observations are discounted in the overall rating at a rate equal to the agent's discount rate. Instead, the characteristics of the signal processes determine the weights of the signal innovations in the incentive state; that is, these characteristics determine the role and relative importance of the signals in the overall rating. Hence, the optimal rating balances the rater's information, as summarized by the rater's belief, with some short-termism that is in proportion to the agent's impatience.[4] Signals that boost career concerns should see their weight

[1] In the case of health care, Dranove, Kessler, McClellan, and Satterthwaite (2003) find that, at least in the short run, report cards decreased patient and social welfare. In the case of education, Chetty, Friedman, and Rockoff (2014a,b) argue that the benefits of value-added measures of performance outweigh the counterproductive behavior that it encourages—but gaming is also widely documented (see Jacob and Lefgren (2005) among many others).

[2] This contrasts with several algorithms based on the principle that the new rating is a function of the old ratings and the most recent review(s) (Jøsang, Ismail, and Boyd (2007)). However, there is also significant evidence that, in many cases, observed ratings (based on proprietary rules) cannot be explained by a simple (time-homogeneous) Markov model. See, among others, Frydman and Schuermann (2008), who precisely argue that two-dimensional Markov models provide a better explanation for actual credit risk dynamics. Such two-state systems are already well-studied under the name of mixture (multinomial) models. See, among others, Adomavicius and Tuzhilin (2005).

[3] In credit ratings, this failure has been widely empirically documented; see Section 3.2.

[4] The ineffectiveness of irrelevant conditioning also resonates with standard principal-agent theory; see, for instance, Green and Stokey (1983).

amplified, while those that stifle career concerns should be muted.

These findings are robust to the informational environment. They hold irrespective of whether past ratings can be hidden from the market (confidential vs. public ratings) and of whether the market has access to additional non-proprietary information (exclusive vs. non-exclusive ratings). However, these distinctions matter for the particulars of the rating mechanism. For instance, if past ratings are observable, then hiding information is only effective if the mechanism has access to diverse sources of information (*i.e.*, multidimensional signals). If it relies on a single source of information, then the best public rating is transparent. Non-exclusivity also matters. In the public case, the mechanism might release more information regarding its hidden sources when others are freely available. Instead, in the confidential case, the free information and that revealed by the rating can be substitutes. In fact, the Markov mixture rating system is canonical, in the sense that any effort that can be induced for some rating system can be induced by a Markov mixture rating system, possibly with some added noise. Hence, their emergence does not rely on our emphasis on effort maximization as objective, or the ability of the rater to commit.

Surprisingly, perhaps, we show that the rating system can count past performance against it. That is, performing well at some point can boost the rating in the short term but depress it in the long term. This is because the impact of a rating is proportional to its scale, the market adjusting for its variance. However, when the agent's ability is not too persistent (low mean-reversion), the variance of the rating is naturally high. By counting recent and older signals in opposing directions, the rating counteracts this. Of course, there is also a direct adverse impact on incentives, but this effect is smaller than the indirect positive effect if the agent is impatient.

Our analysis builds on the seminal model of Holmström (1999).[5] An agent exerts effort unbeknown to the market, which pays him a competitive wage. This wage is based on the market's expectation of the agent's productivity, which depends on instantaneous effort and his ability, a mean-reverting process. This expectation is based on the market's information. Rather than directly observing a noisy signal that reflects ability and effort, the market obtains its information via the rating set by some intermediary. The intermediary potentially has many sources of information about the agent and freely chooses how to convert these signals into the rating. In brief, we view a rating system as an information channel that must be optimally designed. We focus on a simple objective that in our environment is equivalent to social surplus: to maximize the agent's incentive to exert effort or, equivalently, to

[5]Modeling differences with Holmström (1999) include the continuous-time setting, mean-reversion in the type process, and a multidimensional signal structure. See Cisternas (2015) for a specification that is similar to ours in the first two respects.

solve for the range of effort levels that are implementable. (We also examine the trade-off between the level of effort and the precision of the market's information.)[6,7]

We allow for a broad range of mechanisms, imposing that the rating mechanism be time-invariant and its quality be history-independent.[8] As we show, a rating mechanism is equivalent to a time-invariant linear filter, mapping all the bits of information available to the intermediary into a (without loss) scalar rating. In general, such mechanisms are infinite-dimensional.

In Section 4, we study two extensions. First, we allow for ratings that are not exclusive. That is, the market has access to independent public information. We show how the optimal rating reflects the content of this free information. Second, we discuss how our results extend to the case of multiple actions.[9] We show that it can be optimal for the optimal rating system to encourage effort production in dimensions that are unproductive, if this is the only way to also encourage productive effort. Third, we apply our techniques to compare existing methods, showing that exponential smoothing dominates a moving window. An appendix gathers the proofs of Sections 2 and 3. All missing details and proofs can be found in the online Supplementary Appendix.

Related Literature. Foremost, our paper builds on Holmström (1999). (See also Dewatripont, Jewitt, and Tirole (1999).) His model elegantly illustrates why neither perfect monitoring nor a lack of oversight cultivates incentives. His analysis prompts the question raised and answered in our model: what type of feedback

[6] These two objectives feature prominently in economic analyses of ratings according to practitioners and theorists alike. As Gonzalez *et al.* (2004) state, the rationale for ratings stems from their ability to gather and analyze information (information asymmetry) and affect the agents' actions (principal-agent). To quote Portes (2008), "Ratings agencies exist to deal with principal-agent problems and asymmetric information." To be sure, resolving information asymmetries and addressing moral hazard are not the only roles that ratings play. Credit ratings, for instance, play a role in a borrowing firm's default decision (Manso (2013)). Additionally, ratings provide information to the agent himself (*e.g.*, performance appraisal systems); see Hansen (2013). Moreover, whenever evaluating performance requires input from the users, ratings must account for their incentives to experiment and report (Kremer, Mansour, and Perry (2014), Che and Hörner (2015)).

[7] Throughout, we ignore the issues that rating agencies face in terms of possible conflict of interest and their inability to commit, which motivates a broad literature.

[8] As we show, this calls for ratings that are stationary Gaussian processes. Our focus on such mechanisms nonetheless abstracts from some interesting questions, such as the granularity of the rating (the rating's scale) or its periodicity (*e.g.*, yearly vs. quarterly ratings), as well as how ratings should be adjusted to account for the rated firm's age.

[9] With multidimensional product quality, information disclosure on one dimension may encourage firms to reduce their investments in others, harming welfare (Bar-Isaac, Caruana, and Cuñat (2012)).

stimulates effort? Our interest in multifaceted information is similar to Holmstrom and Milgrom (1991), who consider multidimensional effort and output to examine optimal compensation. Their model has neither incomplete information nor career concerns. Our work is also related to the following strands of literature.

Reputation. The eventual disappearance of reputation in standard discounted models (as in Holmström (1999)) motivates the study of reputation effects when players' memory is limited. There are many ways to model such limitations. One is to simply assume that the market can only observe the last K periods (in discrete time), as in Liu and Skrzypacz (2014). This allows reputation to be rebuilt. Even more similar to our work is Ekmekci (2011), who interprets the map from signals to reports as ratings, as we do. His model features an *informed* agent. Ekmekci shows that, absent reputation effects, information censoring cannot improve attainable payoffs. However, if there is an initial probability that the seller is a commitment type that plays a particular strategy every period, then there exists a finite rating system and an equilibrium of the resulting game such that the expected present discounted payoff of the seller is approximately his Stackelberg payoff after every history. As in our paper, Pei (2015) introduces an intermediary in a model with moral hazard and adverse selection. The motivation is very similar to ours, but the modeling and the assumptions differ markedly. In particular, the agent knows his own type, and the intermediary can only choose between disclosing and withholding the signal, while having no ability to distort its content. Furthermore, in Pei (2015), the intermediary is not a mediator in the game-theoretic sense but a strategic player with her own payoff that she maximizes in the Markov perfect equilibrium of the game.

Design of reputation systems. The literature on information systems has explored the design of rating and recommendation mechanisms. See, among others, Dellarocas (2006) for a study of the impact of the frequency of reputation profile updates on cooperation and efficiency in settings with pure moral hazard and noisy ratings. Inspired by the health care market, Glazer and McGuire (2006) show how obfuscation (providing a rating that is an average of different measures) dominates full transparency, because of adverse selection. This literature abstracts from career concerns, the main driver here.

Design of information channels. There is a vast literature in information theory on how to design information channels, and it is impossible to do it justice. Restrictions on the channel's quality are derived from physical rather than strategic considerations (*e.g.*, limited bandwidth). See, among many others, Chu (1972), Ho and Chu (1972) and, more related to economics, Radner (1961). Design under incentive constraints is recently considered by Ely (2015) and Renault, Solan, and Vieille (2015). However, these are models in which information disclosure is distorted because of the incentives of the users of information; the underlying information process is exogenous.

2 The Model

2.1 Exogenous Information

The relationship involves a long-lived agent (he) and a competitive market (it), mediated by an intermediary (she). We first abstract from the intermediary's objective by treating the information transmitted by the intermediary to the market as exogenous. We then turn to optimizing over information structures in Section 3. Time is continuous, indexed by $t \geq 0$, and the horizon is infinite.

There is incomplete information. The agent's ability, or *type*, is $\theta_t \in \mathbf{R}$. We assume that θ_0 has a Gaussian distribution. It is drawn from $\mathcal{N}(0, \gamma^2/2)$. The law of motion of θ is mean-reverting, with increments

$$\mathrm{d}\theta_t = -\theta_t \, \mathrm{d}t + \gamma \, \mathrm{d}Z_{0,t}, \tag{1}$$

where Z_0 is an independent standard Brownian motion (BM), and $\gamma > 0$.[10] The unit rate of mean-reversion is a mere normalization, as is its zero mean.[11] Mean-reversion ensures that the variance of θ remains bounded, independent of the market information, thereby accommodating a large class of information structures.[12] The noise Z_0 ensures that incomplete information persists and that the stationary distribution is nontrivial. The specification of the initial variance precisely ensures that the process is stationary.

The type affects the distribution over output and signals. Specifically, given some real-valued process A_t (the action of the agent), cumulative output $X_t \in \mathbf{R}$ solves

$$\mathrm{d}X_t = (A_t + \theta_t) \, \mathrm{d}t + \sigma_1 \, \mathrm{d}Z_{1,t}, \tag{2}$$

with $X_0 = 0$. Here, Z_1 is an independent standard Brownian motion, and $\sigma_1 > 0$. We allow for but do not require additional signals of ability.[13] We model such sources of

[10] Throughout, when we refer to an independent standard Brownian motion, we mean a standard Brownian motion independent of all the other random variables and random processes of the model.

[11] Going from a mean-reversion rate of 1 to ρ requires the following changes of variables: $t \mapsto \rho t$, $\gamma \mapsto \gamma/\sqrt{\rho}$, $r \mapsto r/\rho$, $(\alpha_k, \beta_k, \sigma_k) \mapsto (\alpha_k/\rho, \beta_k/\rho, \sigma_k/\sqrt{\rho})$. The optimal confidential rating process, as described in Theorem 3.1, remains well-defined and non-degenerate as $\rho \to 0$. On the other hand, the optimal public rating process becomes transparent in that limit.

[12] An alternative approach that we leave unexplored is to allow for some background learning.

[13] In the case of a company, besides earnings, there is a large variety of indicators of performance

information as processes $\{S_{k,t}\}$, $k = 2, \ldots, K$, which are solutions to

$$\mathrm{d}S_{k,t} = (\alpha_k A_t + \beta_k \theta_t)\,\mathrm{d}t + \sigma_k\,\mathrm{d}Z_{k,t}, \tag{3}$$

with $S_{k,0} = 0$. Here, $\alpha_k \in \mathbf{R}$, $\beta_k \geq 0$ (wlog), $\sigma_k > 0$ and Z_k is an independent standard Brownian motion. For convenience, we set $S_1 = X$ (and $\alpha_1 = \beta_1 = 1$), as output also plays the role of a signal. Alongside some initial (for now, exogenous) sigma-algebra $\mathcal{G}_0$, the random variables $\mathbf{S} := \{S_k\}_{k=1}^K$ are the only sources of information. We refer to the corresponding filtration as $\mathcal{G}$, where $\mathcal{G}_t$ is (the usual augmentation of) $\mathcal{G}_0 \vee \sigma\left(\{\mathbf{S}_s\}_{s \leq t}\right)$. This is the information of the intermediary. The agent observes these signals but also knows his own past effort choices. Note that, like the intermediary, the agent learns about his type over time by observing $\mathcal{G}$. On path, his belief coincides with the intermediary's assessment.

The information available to the market at time t is modeled by a sigma-algebra $\mathcal{F}_t \subseteq \mathcal{G}_t$. We do not impose that $\mathcal{F}$ be a filtration, an important point for the sequel. An (agent) strategy is a bounded process A that is progressively measurable with respect to $\mathcal{G}$. Let $\mathcal{A}$ denote the collection of strategies.[14] Neither the market nor the intermediary observe the process A. Instead, the market forms a conjecture about

(profitability, income gearing, liquidity, market capitalization, etc.). In the case of sovereign credit ratings, Moody's and Standard & Poor's list numerous economic, social, and political factors that underlie their rating (Moody's Investor Services (1991), Moody's Investor Services (1995), Standard & Poor's (1994)); similarly, workers are evaluated according to a variety of performance measures, both objective and subjective (see Baker, Gibbons, and Murphy (1994)).

[14]This is intuitive, but heuristic. Formally, as in continuous-time principal-agent models, to avoid circularity problems where actions depend on the process that they define, (2) and (3) are to be interpreted in the weak formulation of stochastic differential equations (SDE), where Z is a BM that generally depends on A. Specifically, signal processes are defined for a reference effort level (say, 0); one defines S_k as the solution to

$$S_{k,t} = \beta_k \int_0^t \theta_s\,\mathrm{d}s + \sigma_k Z_{k,t},$$

and then $\mathcal{G}$ as the natural augmented filtration generated by the processes S_k alongside $\mathcal{G}_0$, with associated probability measure P^0. Thus, the agent actions do not define the signal process itself, which is fixed *ex ante*. Instead, they define the law of the process: given $A \in \mathcal{A}$, define Z_k^A by $Z_{k,t}^A = Z_{k,t} - \frac{\alpha_k}{\sigma_k^2} \int_0^t A_s\,\mathrm{d}s$. By the Girsanov Theorem, there exists a probability measure P^A such that the joint law of $(\theta, Z_1^A, \ldots, Z_K^A)$ under P^A is the same as the joint law of $(\theta, Z_1, \ldots, Z_K)$ under P^0. Given $A \in \mathcal{A}$, the signal S_k satisfies

$$\mathrm{d}S_{k,t} = (\alpha_k A_t + \beta_k \theta_t)\,\mathrm{d}t + \sigma_k\,\mathrm{d}Z_{k,t}^A,$$

with $Z_{k,t}^A$ a BM under P^A. These are the signals that the intermediary observes.

A, denoted $A^* \in \mathcal{A}$, from which a "belief" P^{A^*} is derived. This belief defines the law of motion of the signals and output processes as if the agent exerted effort A_t^* at time t, instead of A_t. Expectations relative to this measure are denoted $\mathbf{E}^*[\cdot]$. This contrasts with the belief P^A of the agent, which captures the actual law of motion. Expectations relative to P^A are denoted $\mathbf{E}[\cdot]$.[15]

We now turn to payoffs. Given a (cumulative) real-valued transfer process to the agent (a continuous, $\mathcal{F}$-adapted process) π, the market retains

$$\int_0^\infty e^{-rt}(\mathrm{d}X_t - \mathrm{d}\pi_t),$$

whereas the agent receives

$$\int_0^\infty e^{-rt}(\mathrm{d}\pi_t - c(A_t)\,\mathrm{d}t).$$

Here, $r > 0$ is the common discount rate.[16] The cost of effort $c(\cdot)$ is twice differentiable, with $c'(0) = 0$, and $c'' > 0$. The transfer π does not matter for efficiency (joint surplus maximization), which demands setting A_t at the constant solution of $c'(A_t) = 1$.

The equilibrium definition has three ingredients. The first is how transfers are set. We assume that the market is competitive and that there is no commitment, in the sense that output-contingent wages are not allowed. Given the market conjecture A^*, it pays a flow transfer $\mathrm{d}\pi_t$ equal to $\mathbf{E}^*[\mathrm{d}X_t \mid \mathcal{F}_t]$ "upfront" (note that this transfer can be negative). Second, the agent chooses his strategy A to maximize his expected payoff. Third, the market has rational expectations, and hence, its belief about A coincides with the optimal strategy. Because our focus will be on equilibria with deterministic effort, we assume throughout that A^* is a deterministic function of time.

Definition 2.1 *Fix an information structure $\mathcal{F}$. An equilibrium is a profile (A, A^*, π), $A, A^* \in \mathcal{A}$, such that:*

---

[15]Formally, we use the star notation when we refer to the law on $(\theta, S_1, \ldots, S_K)$ induced by P^{A^*} (see ft. 14), *i.e.*, the law of motion of the ability and signal processes from the perspective of the market. We use the no-star notation when we refer to the law on $(\theta, S_1, \ldots, S_K)$ induced by P^A, *i.e.*, the law of motion of the agent according to his own belief. The belief of the market is given by the law of the joint process $(\theta, S_1, \ldots, S_K)$, but as the mean ability is all that is payoff-relevant, we abuse language and often call belief the mean ability. The same remark holds for the agent's belief. We drop the star notation for variance and covariance, for which the distinction is irrelevant.

[16]Only the agent's impatience is relevant for equilibrium analysis, and this is how we interpret the parameter r. However, equal discounting is necessary for transfers to be irrelevant for efficiency.

1. *(Zero-profit) For all* t,

$$\pi_t = \int_0^t \mathbf{E}^*[A_s^* + \theta_s \mid \mathcal{F}_s] \, \mathrm{d}s.$$

2. *(Optimal effort)*

$$A \in \operatorname*{argmax}_{\widehat{A} \in \mathcal{A}} \mathbf{E} \left[\int_0^\infty e^{-rt} (\mathrm{d}\pi_t - c(\widehat{A}_t) \, \mathrm{d}t) \right].$$

3. *(Correct beliefs) It holds that*

$$A^* = A.$$

This paper is concerned with the optimal design of the information structure. An important special case is obtained for $\mathcal{F} = \mathcal{G}$ such that the market observes all there is to observe, save for the actual effort. With some abuse, we refer to this case as the model of Holmström, or as *transparency*. However, many more structures are considered. The following are two important properties of information structures.

Definition 2.2 *An information structure $\mathcal{F}_t (\subseteq \mathcal{G}_t)$ is public if $\mathcal{F}$ is a filtration.*

Hence, an information structure is public if all information available to the market in the past remains available at later times. We say that the information structure is *confidential* to insist that we do not require, but do not rule out, that it is public.[17]

Definition 2.3 *An information structure $\mathcal{F}_t (\subseteq \mathcal{G}_t)$ is non-exclusive (w.r.t. signals $K' \subseteq \{1, \dots, K\}$) if*

$$\sigma\left(\{S_{k,t}\}_{k \in K'}\right) \subseteq \mathcal{F}_t.$$

Informally, non-exclusivity means that some signals are observed by the market over time. When such a restriction is not imposed, the information structure is *exclusive* (with further abuse, as non-exclusive information structures are instances of exclusive ones). Non-exclusivity with respect to X is a natural case to consider because this information can be backed out from the payoff process of the market.[18]

Because the payments received by the agent reflect the market belief concerning his type, the agent has incentives to affect this belief via his effort. Hence, given the

[17] Confidentiality can be defined in a stronger sense: different market participants at time t could receive different information. This makes no difference to the analysis.

[18] The relative importance of exclusive vs. non-exclusive information varies substantially across and within industries: in the credit rating industry, solicited ratings are based on both public and confidential information; unsolicited ratings, by contrast, rely exclusively on public information.

equilibrium payment, the agent maximizes his discounted reputation, net of his cost of effort, as formalized below in (4). Fixing the conjecture A^*, a sufficient statistic for $\mathcal{F}$ is the conditional expectation $\mathbf{E}^*[\theta_t \mid \mathcal{F}_t]$. This is all the information that matters for equilibrium analysis.[19] Unless stated otherwise, all proofs are in the appendices.

Lemma 2.4

1. *Given a payment process that satisfies the zero-profit condition, the effort strategy A maximizes the agent's payoff if, and only if, it maximizes*

$$\mathbf{E}\left[\int_0^\infty e^{-rt}(\mu_t - c(A_t))\,\mathrm{d}t\right] \tag{4}$$

over $\mathcal{A}$, where $\mu_t := \mathbf{E}^*[\theta_t \mid \mathcal{F}_t]$ is derived using A^* as the market conjecture.

2. *If $(\mathcal{A}, \pi)$ is an equilibrium given $\mathcal{F}$, then it is an equilibrium given $\{\sigma(\mathbf{E}^*[\theta_t \mid \mathcal{F}_t])\}_t$.*

2.2 Ratings

The intermediary selects an information structure $\mathcal{F}$. *A priori*, such a structure can be arbitrarily complex. However, given Lemma 2.4, the equilibrium effort when the market observes $\mathcal{F}_t$ at time t is identical to the equilibrium effort when the market observes the scalar $\mathbf{E}^*[\theta_t \mid \mathcal{F}_t]$ only. Hence, without loss, it can be assumed that the intermediary releases a scalar rating to the market, Y_t, at time t. Figure 1 summarizes how participants interact.

We focus on stationary environments. This requires defining $\mathcal{G}_0$ such that the environment is *as if* time began at $-\infty$. One way to do so is to regard signals S_k and θ as two-sided processes.[20,21] *Rating processes* are a special class of scalar ratings,

[19] Note that, unlike $\mathcal{F}$, the information structure $\{\mathbf{E}^*[\theta_t \mid \mathcal{F}_t]\}_t$ refers to the market conjecture.

[20] A two-sided process is defined on the entire real line, as opposed to the nonnegative half-line. In particular, we call two-sided standard Brownian motion any process $Z(= \{Z_t\}_{t\in\mathbf{R}})$ such that both $\{Z_t\}_{t\geq 0}$ and $\{Z_{-t}\}_{t\geq 0}$ are standard Brownian motions. Let $\{\mathcal{G}_t\}_{t\in\mathbf{R}}$ be the natural augmented filtration generated by $\{S_k\}_k$, which induces the filtration $\mathcal{G}$ on the nonnegative real line.

[21] Formally, for all $t \in \mathbf{R}$,

$$\theta_t = e^{-t}\bar{\theta} + \int_0^t e^{-(t-s)}\gamma\,\mathrm{d}Z_{0,s},$$

where $\bar{\theta} \sim \mathcal{N}(0, \gamma^2/2)$, and Z_0 is two-sided. Similarly, let $X = S_1$ and, given the two-sided BM Z_k, S_k be the two-sided process defined by (see ft. 14),

$$S_{k,t} = \beta_k \int_0^t \theta_s\,\mathrm{d}s + \sigma_k Z_{k,t}.$$

defined as follows.

Definition 2.5 *A (two-sided) process Y is a rating process if, for all $t \in \mathbf{R}$, Y_t is $\mathcal{G}_t$-measurable, and, when the agent's effort is constant over time,*

(1) for all $\tau > 0$, $(\bar{Y}_t, \mathbf{S}_t - \mathbf{S}_{t-\tau})$ is jointly stationary and Gaussian, where $\bar{Y}_t := Y_t - \mathbf{E}[Y_t]$ is the mean-normalized rating;

(2) for all k, $\tau \mapsto \mathbf{Cov}[Y_t, S_{k,t-\tau}]$ is absolutely continuous, with integrable and square-integrable generalized derivative.

We restrict attention to information structures induced by rating processes. A rating process Y induces an information structure $\mathcal{F}$ via $\mathcal{F}_t = \sigma(Y_t)$. We say that Y is a *confidential/public rating process* when $\mathcal{F}$ is a confidential/public information structure. Unless stated otherwise, Y is not deterministic and we assume throughout $\mathbf{Var}\,[Y_t] > 0$ for all t.

Rating processes preclude some interesting practices. Normality rules out coarse ratings, for instance.[22] Still, it encompasses a variety of rating practices. In the case of a one-dimensional signal, for instance, the process can involve *exponential smoothing* (as allegedly used by Business Week in its business school ranking), which involves setting

$$Y_t = \int_{s \le t} e^{-\delta(t-s)}\, \mathrm{d}X_s,$$

for some choice of $\delta > 0$. The rating system can be a *moving window* (as commonly used in consumer credit ratings or Better Business Bureau (BBB) grades) when

$$Y_t = \int_{t-T}^{t} \mathrm{d}X_s,$$

for some $T > 0$. (In both cases, the choice of $\mathcal{G}_0$ ensures that this is also well defined for $s \le 0$.) A comparison between these two ratings is given in Section 4.3.

Furthermore, it is not difficult to think of special cases in which there exists a scheme that would boost effort beyond what is possible with a rating process. Suppose, for instance, that one of the signals perfectly reveals the agent's effort.[23] Then, it suffices for the rating system to raise a "red flag" (ostensibly ceasing to provide any rating in the future) as soon as it detects a deviation from the desired effort level to ensure that any deviation is unattractive in the first place.

[22] The restriction to adapted processes also rules out the use of extraneous noise by the intermediary. This is merely a modeling choice, as white noise can be included as a dummy signal.

[23] More sophisticated schemes can be devised that apply when there is some small noise in the signal regarding effort while inducing efficient effort at all times.

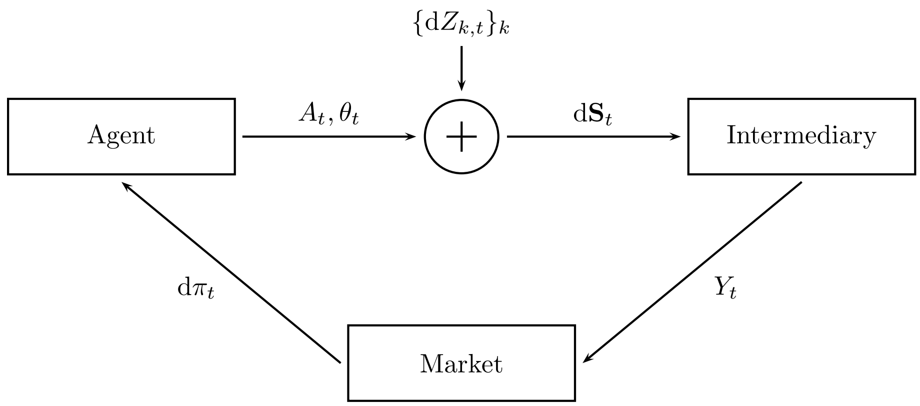

Figure 1: Flow of information and payments between participants.

Yet, we view such schemes as unrealistic for many applications. In punishing the agent, the rating also "punishes" the market by worsening the information it provides. We believe it is both a realistic and desirable property for a rating system that the history affect the content of the rating, but not the quality of the information that it conveys.

Perhaps surprisingly, normality is a *consequence* of assuming that history does not matter for the quality of information: as we prove in the online Supplementary Appendix (see Lemma SA.3 in Section SA.3.2), as long as the covariances (between rating and type, and rating and signal) are deterministic processes (with some smoothness and integrability properties), the non-stationary version of our Representation Lemma (Lemma 2.11) below must hold, and hence the rating *must* have a Gaussian distribution. We further impose stationarity, as the quality of the rating should not depend on calendar time either.

Gaussian processes make the model tractable. It allows to apply linear filtering techniques that we adapt to our framework. Stationarity ensures that equilibrium effort is a scalar, facilitating comparisons. From now on, A^* is taken to be constant on $\mathbf{R}_+$ and the equilibrium to be stationary (set by convention $A_s = A^*_s = 0$ if $s < 0$).

Rating processes can represent beliefs for both confidential and public information structures. They admit the following immediate and intuitive characterization.

Proposition 2.6 *Let Y be a rating process. Then, Y is as follows:*

1. *A belief for a confidential information structure if, and only if, for all* t,

$$\mathbf{E}^*[\theta_t \mid Y_t] = Y_t.$$

2. *A belief for a public information structure if, and only if, for all* t,

$$\mathbf{E}^*[\theta_t \mid \{Y_s\}_{s \leq t}] = Y_t.$$

The following provides a simple criterion to decide whether a rating process is equal to a market belief of a confidential or a public information structure.[24]

Lemma 2.7 (Confidential Belief) *A rating process* Y *is a belief for a confidential information structure if, and only if, for all* t,

$$\mathbf{E}^*[Y_t] = 0 \textit{ and } \mathbf{Cov}[Y_t, \theta_t] = \mathbf{Var}[Y_t].$$

Hence, the lemma implies that any rating process with mean zero is proportional to the mean belief that it induces.

Lemma 2.8 (Public Belief) *A rating process* Y *is a belief for a public information structure if, and only if, it is a belief for a confidential information structure and in addition, for all* t *and all* $\tau \geq 0$,

$$\mathbf{Corr}[Y_t, Y_{t+\tau}] = \mathbf{Corr}[\theta_t, \theta_{t+\tau}] \ (= e^{-\tau}).$$

Instead of focusing on beliefs, it is often convenient to work with a slighter broader class of rating processes. Scaling a rating process by a nonzero constant does not affect its informational content. Hence, we may select as convenient a rating process within the equivalence class that this constant of proportionality defines.

[24] The characterization is helpful to compute the optimal ratings, as it allows us to restrict attention to ratings that are belief processes. The optimization is then performed under a set of constraints that we relax by internalizing them in the objective function.

As an alternative approach, we could optimize over the general, unconstrained family of rating processes Y, and compute the associated beliefs $\mathbf{E}[\theta_t \mid Y_t]$ or $\mathbf{E}[\theta_t \mid \{Y_s\}_{s \leq t}]$. This can be done for confidential information structures (with or without exclusive information). However, for public information structures, the computation of beliefs is not tractable. In particular, since Y is not required to have a Markovian structure, common linear filters such as the Kalman-Bucy filter cannot be used. Instead, the computation of beliefs involves the determination of a continuum of variables associated with conditional variances of ratings that solve a continuum of equations, the analytic solution to which can only be written in some special cases.

2.3 Characterization of Equilibrium

Lemma 2.9 *Fix a rating process Y. Under the information structure it induces, $\mathcal{F} = \{\sigma(Y_t)\}_{t\geq 0}$, there exists a unique equilibrium.*

We have assumed that the agent observes the signal realizations over time, as captured by the filtration $\mathcal{G}$. In turn, such information defines the set of agent strategies. But one may also want to consider an environment in which the agent observes different information, modeled by some arbitrary filtration $\mathcal{H}$, allowing for another set of strategies.[25] The concept of equilibrium, presented in Definition 2.1, depends on the agent information only through the set of agent strategies, and so remains valid without modification. The same observation applies to the proof of Lemma 2.9, which does not depend on the specific information available to the agent, because the optimal effort level specified by the equilibrium, turns out to be constant. Therefore, the following corollary immediately obtains.

Corollary 2.10 *Fix an arbitrary filtration $\mathcal{H}$. If the agent instead observes information captured by $\mathcal{H}$ (as opposed to $\mathcal{G}$) there continues to exist a unique equilibrium, and the equilibrium effort level does not depend on the chosen filtration $\mathcal{H}$.*

We now turn to the characterization. This is done in two stages. First, the restriction to rating processes leads to a convenient analytic representation.

Lemma 2.11 (Representation Lemma) *Fix a rating process Y. Given any market conjecture A^*, there exist essentially unique*[26] *integrable and square-integrable functions u_k, $k = 1, \ldots, K$, such that*

$$Y_t = \mathbf{E}^*[Y_t] + \sum_{k=1}^{K} \int_{s\leq t} u_k(t-s)(\mathrm{d}S_{k,s} - \alpha_k A_s^* \,\mathrm{d}s). \tag{5}$$

The coefficient $u_k(s)$ is the weight that the current rating Y_t attaches to the innovation (the term $(\mathrm{d}S_{k,s} - \alpha_k A_s^* \,\mathrm{d}s)$) pertaining to the signal of type k and vintage s. Following information-theoretic terminology, we refer to $\{u_k\}_k$ as the *linear filter* defined by Y. When the filter is a sum of exponentials (*e.g.*, $u_k(t) = \sum_\ell c_\ell e^{-\delta_\ell t}$), the coefficients (resp., exponents) are the *weights* (resp., *impulse responses*) of the filter. Conversely, given some filter $\{u_k\}_k$, (5) uniquely defines a rating process.

[25] Formally, an agent strategy is then a bounded process progressively measurable with respect to $\mathcal{H}$.

[26] Unique up to measure zero sets.

The decomposition of Lemma 5 can be interpreted as a regression of Y_t on the continuum of signal increments $\mathrm{d}S_{k,s}$, $s \in (-\infty, t]$.[27] It is an infinite-dimensional version of the familiar result that a Gaussian variable that is a dependent function of finitely many Gaussian variables is a linear combination thereof.[28] It turns the optimization problem into a deterministic one. There is an explicit formula for u_k in terms of the covariance of the rating, which shows how u_k captures not only the covariance between the rating and a weighted average of the signals of a given vintage but also how this covariance decays over time for signal k. For all $t \geq 0$,

$$u_k(t) = \frac{\beta_k \gamma^2}{\sigma_k^2 \kappa} \left(\frac{\sinh \kappa t + \kappa \cosh \kappa t}{1+\kappa} \int_0^\infty e^{-\kappa s} \, \mathrm{d}\bar{f}(s) - \int_0^t \sinh \kappa(t-s) \, \mathrm{d}\bar{f}(s) \right) - \frac{f_k'(t)}{\sigma_k^2},$$

with $\kappa := \sqrt{1 + \gamma^2 \sum_k \beta_k^2/\sigma_k^2}$ (> 1), and

$$f_k(s) := \mathbf{Cov}[Y_t, S_{k,t-s}], \text{ and } \bar{f}(s) := \sum_{k=1}^K \frac{\beta_k}{\sigma_k^2} f_k(s).$$

Second, we express equilibrium effort in terms of the filter given by Lemma 5.

Lemma 2.12 *Let Y be a rating process with normalized variance, $\mathbf{Var}[Y_t] = 1$.*[29] *The unique equilibrium effort level A is constant and determined by*

$$c'(A) = \frac{\gamma^2}{2} \left[\sum_{k=1}^K \alpha_k \int_0^\infty u_k(t) e^{-rt} \, \mathrm{d}t \right] \left[\sum_{k=1}^K \beta_k \int_0^\infty u_k(t) e^{-t} \, \mathrm{d}t \right], \tag{6}$$

where u_k is defined by Lemma 2.11, given Y.

Hence, effort is proportional to the product of two covariances. The first pertains to the agent: the impact of effort and his discount rate. The other pertains to the type: the impact of ability and the mean-reversion rate. This formula assumes a

[27] Determining the coefficients of such continuous-time regressions is often achieved via a linear filtering argument. Here, the lack of Markovian structure with the infinite fictitious history, together with the stationarity condition, makes the problem non-trivial because it prevents the use of the Kalman-Bucy filter and involves finding a continuum of terms of the form $\mathbf{Var}[Y_t \mid \mathcal{G}_{t-s}]$ that solve a continuum of equations. To obtain the closed-form solution for the coefficients u_k, we write the equations that link f to u_k; then, via algebraic manipulation and successive differentiation, we obtain a differential equation that u_k must satisfy, the solution of which is found explicitly.

[28] Given the continuum, stronger assumptions are necessary. The restriction to stationary processes is key to obtaining a linear representation. See Jeulin and Yor (1979) for a counter-example otherwise.

[29] The somewhat unwieldy statement of this constraint in terms of $\{u_k\}_k$ is given in (15) below.

normalized variance. Alternatively, we may write (6) in a compact way as

$$c'(A^*) = \left[\sum_{k=1}^{K} \alpha_k \int_0^\infty u_k(t) e^{-rt} \, \mathrm{d}t\right] \frac{\mathbf{Cov}[Y_t, \theta_t]}{\mathbf{Var}[Y_t]}. \tag{7}$$

The objective of Section 3 is to find the rating process that maximizes the right-hand side of (6), under the constraints imposed by Lemmas 2.7 and 2.8.

2.4 Transparency

Here, we consider the benchmark in which $\mathcal{F} = \mathcal{G}$. This case is close to the one considered by Holmström (1999) in discrete time (specifically, his Section 2.2 solves for the stationary equilibrium with one signal, and no mean-reversion). Define:

$$m_\alpha = \sum_{k=1}^{K} \frac{\alpha_k^2}{\sigma_k^2}, \qquad m_{\alpha\beta} = \sum_{k=1}^{K} \frac{\alpha_k \beta_k}{\sigma_k^2}, \qquad m_\beta = \sum_{k=1}^{K} \frac{\beta_k^2}{\sigma_k^2}. \tag{8}$$

The belief of the market μ_t (defined as $\mathbf{E}^*[\theta_t \mid \mathcal{F}_t]$) is then equal to the intermediary's belief $\nu_t \coloneqq \mathbf{E}^*[\theta_t \mid \mathcal{G}_t]$. The latter is a Markov process that solves

$$\mathrm{d}\nu_t = (\kappa - 1) \sum_k \frac{1}{m_\beta} \frac{\beta_k}{\sigma_k^2} (\mathrm{d}S_{k,t} - \alpha_k A_t^* \, \mathrm{d}t) - \kappa \nu_t \, \mathrm{d}t, \tag{9}$$

where $A^* \in \mathbf{R}_+$ is equilibrium effort. Explicitly, the belief is equal to

$$\nu_t = (\kappa - 1) \int_{s \le t} e^{-\kappa(t-s)} \sum_k \frac{1}{m_\beta} \frac{\beta_k}{\sigma_k^2} \left(\mathrm{d}S_{k,s} - \alpha_k A_s^* \, \mathrm{d}s\right). \tag{10}$$

Innovations $(\mathrm{d}S_{k,s} - \alpha_k A^* \, \mathrm{d}s)$ are weighted according to their type and their vintage. A signal of type k is weighted by the signal-to-noise ratio β_k/σ_k^2.[30] If it is noisy (high σ) or insensitive to ability (low β), it matters little for inferences. Given that ability changes, older signals matter less than recent ones: a signal of vintage $t - s$ is discounted (in the belief at time t) at rate κ. The market "rationally forgets."

Theorem 2.13 *The unique equilibrium effort level when $\mathcal{F} = \mathcal{G}$ is the solution to*

$$c'(A) = \frac{1}{\kappa + r} \frac{m_{\alpha\beta}}{m_\beta} (\kappa - 1), \tag{11}$$

[30]Specifically, β_k/σ_k^2 is the inverse of the Fano factor (the signal-to-noise ratio is β_k/σ_k).

if the right-hand side of (11) *is nonnegative. Otherwise,* $A = 0$.

Equation (11) is a standard optimality condition for investment in productive capital. The market's belief is an asset. Effort is an investment in that asset. We interpret the three terms in (11) as persistence ($(\kappa + r)^{-1}$), substitutability ($m_{\alpha\beta}$), and sensitivity ($\kappa - 1$). The asset depreciates at rate κ, to be added to the discount rate when evaluating the net present value of effort. Investment has productivity $m_{\alpha\beta}/m_{\beta}$, which measures the increase in belief given a (permanent) unit increase in effort. In turn, sensitivity measures the increase in belief given a unit increase in the type.

Substitutability, sensitivity and persistence already appear in (9). Sensitivity is the first coefficient, scaling the impact of a surprise in the signal on the belief; substitutability appears in the sum, as the impact of effort on the surprise; and persistence enters via the last term, capturing the rate of decay of the belief. Only discounting is missing. The general formula given by (6) shows that, for an arbitrary rating process, effort depends also on a fourth term, the ratio $\mathbf{Cov}[Y, \theta]/\mathbf{Var}[Y]$, which is equal to one under transparency. Persistence, sensitivity and substitutability matter as well, and are all nested in the first term, $\int_{t \geq 0} (\sum_k \alpha_k u_k(t)) e^{-rt} \, dt$.

Effort can be too high or low, according to how (11) compares to one. If $m_{\alpha\beta} < 0$, the agent has perverse career concerns: to impress the market, lower effort is better. As a result, equilibrium effort is 0. If $m_{\alpha\beta} = 0$, effort has no impact on the market belief, and equilibrium effort is also 0. Hereafter, we assume that $m_{\alpha\beta} > 0$. A signal for which $\alpha_k = 0$ is not irrelevant, as it enters both sensitivity and persistence. With no signal beside output, effort is inefficiently low, even as discounting vanishes. This is in contrast to Holmström and is due to mean-reversion, which eventually erases the benefits from an instantaneous effort increase.[31] The proof of the following is immediate and omitted.

Lemma 2.14 *Effort increases in* γ *and* α_k, $k = 1, \ldots, K$. *It decreases in* σ_k *if signals are homogeneous* ($\alpha_k, \beta_k, \sigma_k$ *independent of* k)*. It admits an interior maximum with respect to* β_k.

The comparative statics with respect to α_k, γ need no explanation.[32] The role of β_k is more interesting. If it is small, then the market dismisses signal k in terms of learning. If it is high, then the small variation in the signal caused by an effort increase is (wrongly) attributed to the type, but by an amount proportional to β_k^{-1}, which is small and hence not worth the effort cost: a higher β_k makes signal k more relevant but less manipulable. The "best" signals are those involving intermediate

[31] See Cisternas (2012) for the same observation in a model with human capital accumulation.

[32] Effort need not decrease in σ_k if signals are heterogeneous.

values of β_k. Adding a signal has ambiguous consequences for effort, as should be clear. Depending on parameters, it might reduce noise, and so bolster incentives, but it might help tell apart effort and ability, and so undermine career concerns. Either way, it improves the quality of the market's information.

2.5 The Role of the Intermediary

The intermediary's objective is to maximize equilibrium effort A. Her instrument is information. Controlling directly the information available to the agent (an instrument considered by, for instance, Hansen (2013)) is useless: by Corollary 2.10, the equilibrium effort level is independent on the information provided to the agent.

On the other hand, the information provided to the market has an indirect impact on effort. Hence, we focus on her choice of the (market) information structure $\mathcal{F}$ via a rating process Y, which, by our earlier results, we can consider scalar and proportional to the market mean belief it induces. Recall from Lemma 2.9 that the stationary equilibrium is unique, and hence that her choice of Y determines A. She might face constraints: the information structure that the rating defines might be public, non-exclusive, or both. She has commitment, in the sense that Y is chosen once and for all, and it is common knowledge.[33]

Maximizing A does not always maximize efficiency. Even under transparency, equilibrium effort can be too high (cf. (11)). However, solving for the maximum effort is equivalent to solving for the range of implementable actions. If effort is excessive, a simple adjustment to the rating process (adding "white noise," for instance) scales it to any desired lower effort, including the efficient level.

Lemma 2.15 *Fix a confidential/public rating process Y such that stationary effort is A.*

For all $A' \in [0, A]$, there exists a confidential/public rating process Y' such that, under the information structure defined by Y', equilibrium effort is A'.

However, under non-exclusivity, there can be a strictly positive lower bound on the effort that the intermediary can induce. (The information structure that Y' defines in Lemma 2.15 might violate non-exclusivity constraints satisfied by Y.) Surprisingly, this lower bound is not typically achieved by silence (the intermediary disclosing no

[33] This intermediary can be regarded as a "reputational intermediary," an independent professional whose purpose is to transmit a credible quality signal about the agent. Commitment, then, results from the professional's incentive to preserve his reputation. Reputational intermediaries not only include so-called rating agencies but also, in some of their roles, underwriters, accountants, lawyers, attorneys and investment bankers (see Coffee (1997)).

information). Maximum and minimum effort are dual problems. In the presence of non-exclusivity, an optimized information structure can depress effort below what silence achieves, just as it can motivate effort beyond what transparency achieves. See Section 4.1 for further discussion.

Hence, our goal is primarily normative: to identify the range of implementable actions. Yet, there are plausible scenarios in which a profit-maximizing rating agency would find it optimal to induce the maximum effort level. For instance, if the agency charges the market a commission (a set percentage of the value of output), then maximizing effort is equivalent to maximizing revenue.

Depending on the context, it might be desirable to evaluate the performance of a rating process along other dimensions, for instance, the quality of the information it conveys (as measured by the variance of the type conditional on the belief), or its stability over time (as measured by the variance of the belief).[34] These properties satisfy a simple relationship.

Lemma 2.16 *Fix a rating process Y. It holds that*

$$\mathbf{Var}[\theta_t \mid \mu_t] + \mathbf{Var}[\mu_t] = \frac{\gamma^2}{2} \ (= \mathbf{Var}[\theta_t]).$$

Hence, precision and stability are perfect substitutes. If stability comes first, lower precision is desirable. This also means that we can restrict attention to one of these measures when evaluating the trade-off with maximum effort. A systematic analysis of this trade-off would take us too far afield, but it can be done, as illustrated in Section 3.5.

3 Main Results

3.1 Persistence vs. Sensitivity: Two Examples

To build intuition, let us begin with a simple example: exponential smoothing as a confidential rating. Suppose that the intermediary wishes to use the rating

$$Y_t = \sum_k \frac{\beta_k}{\sigma_k^2} \int_{s\le t} e^{-\delta\kappa(t-s)}\, \mathrm{d}S_{k,s},$$

where she freely selects $\delta > 0$. The choice $\delta = 1$ reveals her own belief, and transparency results, as in Section 2.4. Any other choice of δ implies that the market

[34]These properties of ratings are often cited as being desirable (Cantor and Mann (2007)).

is less well-informed than she is. Using the formula from Lemma 2.12, we obtain

$$c'(A) = \frac{1}{\delta\kappa + r} \frac{m_{\alpha\beta}}{m_\beta} \frac{\delta(\kappa+1)(\kappa-1)}{\kappa+\delta}.$$

In terms of the effects introduced before, the first factor ($\frac{1}{\delta\kappa+r}$) is persistence. Future returns on effort are discounted both because of impatience and because future ratings discount past signals at rate $\delta\kappa$. Rating persistence increases the impact of current effort on future ratings. However, increasing persistence decreases sensitivity. This is clear from the last term, which increases in δ and goes to zero if δ does. If δ is small, then the rating is very persistent, which means that it treats old and recent innovations symmetrically. Because ability changes over time, this blunts the impact of a one-time innovation in the belief. If, instead, δ is large, the rating disproportionately reacts to recent innovations, heightening their relative importance.

What goes up must come down: in a stationary system, a blip in a signal cannot simultaneously jolt the belief and have its impact linger. The intermediary must trade off persistence with sensitivity. But she can do better than transparency. Taking derivatives (with respect to δ) yields as optimal solution

$$\delta = \sqrt{r}.$$

She chooses a rating process that is more or less persistent than Bayesian updating according to $r \lesseqgtr 1$, that is, depending on how the discount rate compares to the rate of mean-reversion. The best choice reflects agent preferences, which Bayes' rule ignores. If the agent is patient, it pays to amplify persistence, and δ is low.

Let us turn to a richer example. Departing from our convention regarding output, assume that output is solely a function of ability, not of effort ($\beta := \beta_1 > 0, \alpha_1 = 0$), while the unique other signal purely concerns effort ($\alpha := \alpha_2 > 0, \beta_2 = 0$), and set $\sigma := \sigma_1 = \sigma_2$.[35] Consider the best rating system within the two-parameter family

$$u_1(t) = \frac{\beta}{\sigma^2} e^{-\kappa t}, \quad u_2(t) = d \frac{\beta}{\sigma^2} \sqrt{\delta} e^{-\delta t},$$

with $d \in \mathbf{R}, \delta > 0$. This family is special yet intuitive: because the agent cannot affect output, the intermediary does not distort the corresponding innovations. However, she adds to the resulting integral an integral over the innovations of the second signal. The parameter d scales the weight on this term and δ is its impulse response. The

[35]In this example, efficiency requires $A = 0$: an efficient rating process should discourage effort, a trivial endeavor. We seek the effort-maximizing scheme.

normalization constant $\sqrt{\delta}$ ensures that the choice of δ does not affect the variance of the market belief.[36] Using Lemma 2.12 here as well,

$$c'(A) = \frac{1}{\delta + r}\sqrt{\delta} m_{\alpha\beta} d \frac{2}{(1+d^2)(1+\kappa)}. \tag{12}$$

The first term is familiar by now: it is the impact on persistence of the choice of δ. An effort increase at time t is reflected in the rating at time $t + \tau > t$ but discounted twice: at a rate $e^{-r\tau}$ by the agent and $e^{-\delta\tau}$ by the market. Integrating over $\tau \geq 0$ yields a boost to incentives proportional to $1/(r + \delta)$, which is further amplified by the factor $\sqrt{\delta}$ that scales substitutability. The constant d increases substitutability. But increasing it also increases the belief variance, depressing sensitivity and, hence, effort. This is reflected by the denominator $1 + d^2$. If d is too small, sensitivity disappears, as δ is useless if the second signal does not enter the rating; if too large, sensitivity vanishes because ratings no longer inform ability. An intermediate value is best. The maximization problem is separable, as is clear from (12); $d/(1 + d^2)$ is maximized at $d = 1$ and $\sqrt{\delta}/(r + \delta)$ at $\delta = r$. Independent of δ, the optimal weight on the second term is 1; independent of d, the choice of impulse response is r.

Hence, the intermediary does not only wish to influence persistence, by distorting via δ the weights assigned to signals of different vintages. Via d, she also manipulates the weights assigned to signals of different types to influence substitutability.

3.2 Optimal Ratings

This section solves for the optimal exclusive rating processes. We assume throughout that κ, κ^2, r, and 1 are all distinct. Define

$$\lambda = (\kappa - 1)\sqrt{r}(1 + r)m_{\alpha\beta} + (\kappa - r)\sqrt{\Delta}, \quad \Delta = (r + \kappa)^2(m_\alpha m_\beta - m_{\alpha\beta}^2) + (1 + r)^2 m_{\alpha\beta}^2.$$

Theorem 3.1 *The optimal confidential rating process is unique and given by*[37,38]

$$u_k^c(t) = \frac{\beta_k}{\sigma_k^2}\left(d_k^c \frac{\sqrt{r}}{\lambda} e^{-rt} + e^{-\kappa t}\right),$$

[36] Plainly, once u_2 is squared and integrated over all $\tau \geq 0$, δ vanishes.

[37] Recall that we take ratings as proportional to the market mean belief. Throughout, uniqueness is to be understood as up to such a transformation.

[38] For convenience, the formula here also assumes that $\lambda \neq 0$. The proof gives the general formula.

with coefficients

$$d_k^c := (\kappa^2 - r^2) m_\beta \frac{\alpha_k}{\beta_k} - (\kappa^2 - 1) m_{\alpha\beta}.$$

Theorem 3.2 *The optimal public rating process is unique and given by*

$$u_k^p(t) = \frac{\beta_k}{\sigma_k^2} \left(d_k^p \frac{\sqrt{r}}{\lambda} e^{-\sqrt{r}t} + e^{-\kappa t} \right),$$

with coefficients

$$d_k^p := \frac{\kappa - \sqrt{r}}{\kappa - r} d_k^c + \lambda \frac{\sqrt{r} - 1}{\kappa - r}. \tag{13}$$

Theorems 3.1 and 3.2 provide solutions that are remarkably similar. Because the linear filter is the sum of two exponentials, the rating can be written as a sum of two Markov processes. That is, in both cases, for some $\phi \in \mathbf{R}$,

$$Y_t = \phi I_t + (1 - \phi) \nu_t,$$

where

$$\mathrm{d}I_t = \frac{\beta_k}{\sigma_k^2} \frac{\sqrt{r}}{\lambda} \sum_k d_k (\mathrm{d}S_{k,t} - \alpha_k A_t^* \,\mathrm{d}t) - \delta I_t \,\mathrm{d}t,$$

with $(d, \delta) = (d^c, r)$ in the confidential case and $(d, \delta) = (d^p, \sqrt{r})$ in the public one. The intermediary combines her own belief ν_t with another Markov process, which we denote I (for "incentive"). Its impulse response reflects the agent's patience, as in the second example of Section 3.1. If he is patient, the rating is persistent. If not, performance is reflected in the rating more rapidly than under Bayes' rule.

This common representation has several consequences:

- The optimal rating is not a Markov process. This echoes a large empirical literature documenting that (bond and credit) ratings do not appear to satisfy the Markov property (Altman and Kao (1992), Altman (1998), Nickell, Perraudin, and Varotto (2000), Bangia *et al.* (2002), Lando and Skødeberg (2002), etc.).
- The optimal rating is not a function of the intermediary's belief alone.[39] At first glance, this might be surprising, as the intermediary's belief is the only payoff-relevant variable (in the confidential case). Hence, one might expect the solution to be Markov with respect to this belief. However, the effort-maximizing way of distributing the

[39]This property is distinct from the first. The rating can be Markov without being a function of the belief (this occurs in the first example of Section 3.1). The rating can be a function of ν without being Markov, as functions of Markov processes typically fail to inherit the Markov property.

impact of an innovation over future ratings is not measurable with respect to the intermediary's beliefs at those future times.

- The optimal rating is a two-state mixture Markov rating—a combination of Markov chains moving at different speeds (Frydman (2005)). Using an EM algorithm, Frydman and Schuermann (2008) find that not only does such a two-state mixture Markov model outperform the Markov model in explaining credit ratings, but it also explains economic features of rating data.

What is most surprising is not that two Markov processes are needed to compute the rating, but that two suffice. The part of the proof establishing sufficiency, explained in Section 3.3, sheds light on this. When regarded as a principal-agent model (the principal is the intermediary), promised utility does not suffice as a state variable. Utility is meted out via the market's belief, and beliefs are correct on average. This imposes a constraint on an auxiliary variable and hence demands a second state.

The incentive state is an abstract construct. Another way of understanding what the intermediary does involves re-writing the rating in terms of another pair of states. Unsurprisingly, there is considerable leeway here. For instance, using (Y, ν) (the rating itself and the intermediary's beliefs) leads to a more concrete if less elegant prescription. Explicitly, in the confidential case,

$$\mathrm{d}Y_t = \frac{\beta_k}{\sigma_k^2}\frac{\sqrt{r}}{\lambda}\sum_k (d_k^c + 1)\left(\mathrm{d}S_{k,t} - \alpha_k A_t^* \,\mathrm{d}t\right) - rY_t \,\mathrm{d}t + \frac{(\kappa+1)(r-\kappa)}{\gamma^2}\nu_t \,\mathrm{d}t,$$

and hence, the intermediary continues to incorporate some of her private information (via her belief ν_t) into the rating. In terms of (Y, ν), Y is a hidden Markov process, with ν as the hidden state. This is the formulation occasionally considered for empirical purposes; see Giampieri, Davis, and Crowder (2005). Other representations of the rating process are possible, of course (*e.g.*, as a process with rating momentum; see Stefanescu, Tunaru, and Turnbull (2006)).

Consider now the specific coefficients of the optimal rating processes. The following holds for both the optimal confidential and public ratings. *White noise is harmful*: if $\alpha_k = \beta_k = 0$, then signal k's weight in the rating is zero. Irrelevant noise has no use, as it depresses effort. *All signals enter the rating*: except for a non-generic parameter set, the rating involves them all. Some might be weighted negatively, when innovations along that dimension adversely impact incentives. However, as long as a signal is informative of at least type or effort, the rating takes it into account.

Given the proof of Lemma 2.15, and given Theorems 3.1–3.2, it follows that any effort level $A' \in [0, A]$ (where A is maximum equilibrium effort) can be achieved by considering a rating that is an average of Y_t, the rating as defined by the theorems, and pure "white noise" W_t, an arbitrary independent Wiener process. Hence, two-state Markov mixture ratings are canonical, in the sense that they are optimal, independently of the exact objective of the intermediary, its commitment ability and market power.

Corollary 3.3 *Any stationary public/confidential equilibrium effort level can be achieved by a two-state mixture Markov model plus white noise.*

Among the differences between public and confidential ratings, two are notable. First, the impulse response on the incentive state conspicuously differs across the two environments: this state decays at the discount rate r in the confidential case, whereas in the public case, it does so at a rate equal to the geometric mean between discounting and mean-reversion: $\sqrt{r} = r^{\frac{1}{2}}1^{\frac{1}{2}}$ (the correct interpretation, using the change of variables in ft. 11). The reason is simple. As the second example in Section 3.1 suggests, the impulse response that best trades off persistence with sensitivity is r. Unfortunately, the resulting autocorrelation fails to align with that of a public belief, which decays at the mean-reversion rate (Lemma 2.8). The optimal public rating fixes this in two ways: it distorts the impulse response on the incentive state away from the discount rate toward the mean-reversion rate, and it skews the weight on the incentive term d_k^p away from its favorite weight d_k^c (see (13)).

The second difference is concealed in the definition of this weight d_k^p. If signals are identical (more generally, if, and only if, the ratio α_k/β_k is the same for all signals that are not white noise), then these weights are all zero, and transparency is obtained. While this condition is non-generic for $K > 1$, it is always true when output is the only signal. Instead, with confidential ratings, transparency is a non-generic phenomenon, independent of K. The problem with one signal only is that twisting a weight and an impulse response partially is insufficient to fix the autocorrelation. The weight must be taken all the way down to zero: the "continuum" of autocorrelation constraints determines the "one-dimensional continuum" of variables (the filter $u_1(\cdot)$), and hence the rating, up to some white noise that the intermediary does not wish to use.

To conclude this section, we note that the weights of different signals in the incentive term are ordered according to α_k/β_k.[40]

[40]However, whether the ranking increases or decreases in the ratio depends not only on the sign of m_β but also on whether $r \in [\kappa, \kappa^2]$ ($r \leq \kappa^2$) in the confidential (public) case.

3.3 Proof Overview

Our problem has some unconventional features, so that applying dynamic programming or Pontryagin's maximum principle directly (as is usually done in principal-agent models) is difficult. Hence, our method of proof is somewhat non-standard. Hopefully, it may be useful in related contexts.

In the first part (available in the online Supplementary Appendix), we derive necessary conditions using calculus of variations. These conditions determine a unique candidate for the optimal rating (up to a factor), if it exists and is sufficiently regular. In the second part, we verify that the guess from the first part is optimal. This step introduces a parameterized family of auxiliary principal-agent models and takes limits in a certain way.

Part I: Necessary Conditions

Recall that the ratings communicated to the market may be confidential or public, and the information generated by the signals exclusive or non-exclusive. Thus, there are four settings of interest. In all settings, we normalize the mean rating to zero, and the variance to one.

The Representation Lemma (Lemma 2.11) characterizes all rating processes in terms of a linear filter $\mathbf{u}$, which we use as a control variable. Lemma 2.12 expresses the equilibrium marginal cost of the agent as a function of the filter. Maximizing the equilibrium action is equivalent to maximizing the marginal cost. Thus, we seek to identify a control $\mathbf{u}$ that maximizes a product of two integrals over $\mathbf{u}$:

$$\frac{\gamma^2}{2}\left[\sum_{k=1}^{K}\alpha_k\int_0^{\infty}u_k(t)e^{-rt}\,\mathrm{d}t\right]\left[\sum_{k=1}^{K}\beta_k\int_0^{\infty}u_k(t)e^{-t}\,\mathrm{d}t\right]. \tag{14}$$

In this first part of the proof, we focus on controls that exhibit a sufficient degree of regularity, and we assume that a solution exists within that family.

The maximization is subject to the constraints that the rating process must satisfy. In the simplest case of confidential exclusive information structures, the only constraint is the variance normalization, which is written as follows:

$$\sum_{k=1}^{K}\sigma_k^2\int_0^{\infty}u_k(s)^2\,\mathrm{d}s+\frac{\gamma^2}{2}\int_0^{\infty}\int_0^{\infty}U(s)U(t)e^{-|s-t|}\,\mathrm{d}s\,\mathrm{d}t=1, \tag{15}$$

where $U := \sum_k \beta_k u_k$. The higher dimensionality of the problem is plain in (15). Maximizing (14) subject to (15) is a variational problem with an isoperimetric constraint. We form the Lagrangian and consider a relaxed, unconstrained problem that "internalizes" the variance normalization as part of the objective function. However, the problem is not standard: both objective (14) and constraint (15) include multiple integrals, yet the control has a one-dimensional input. Adapting standard arguments, we prove a version of the Euler-Lagrange necessary condition that covers our class of programs (see Section SA.3.1 of the online Supplementary Appendix). This condition takes the form of an integral equation in $\mathbf{u}$, which can be solved in closed form via successive differentiation and algebraic manipulation. The solution of the relaxed problem can be shown to be a solution of the original problem, which yields a candidate for the optimal rating (unique subject to regularity conditions).

In the more general public and/or non-exclusive settings (see Section 4.1), the objective (14) remains the same, but there are additional constraints on the rating process. These capture the restriction that market beliefs are linked to public or non-exclusive information structures. Lemmas 2.7 and 2.8 state these constraints in the exclusive case, and Lemma 4.2 does so for the non-exclusive cases. Then, we can apply the Representation Lemma (Lemma 2.11) to express these constraints in terms of the filter $\mathbf{u}$ directly.

There are two additional difficulties in these settings. First, there is no longer a finite number of constraints, but a continuum of them. Second, these constraints involve further integral equations with delay.[41] For example, in the public exclusive setting, the constraint (15) is replaced by

$$\sum_{k=1}^{K} \sigma_k^2 \int_0^\infty u_k(t) u_k(t+\tau)\,\mathrm{d}t + \frac{\gamma^2}{2}\int_0^\infty \int_0^\infty U(s)U(t)e^{-|s+\tau-t|}\,\mathrm{d}s\,\mathrm{d}t = 1, \qquad \forall \tau \geq 0. \quad (16)$$

To address this, we reduce the continuum of constraints to a finite set of constraints, applying "educated" linear combinations. We solve the relaxed optimization problem with a finite number of constraints in a manner similar to that for the simplest setting

[41]There is a small literature on the calculus of variations with delayed arguments for single integrals. See Kamenskii (2007) and references therein. There is also a literature on multiple integrals without delayed arguments; see Morrey (1966) for a classical treatise. In both cases, the domain of the control is of the same dimension as the domain of integration.

just described. For instance, in the public exclusive setting, we replace (16) by

$$
\begin{aligned}
1 &= \sum_{k=1}^{K} \sigma_k^2 \int_0^\infty u_k(t)u_k(t)\,\mathrm{d}t + \frac{\gamma^2}{2}\int_0^\infty \int_0^\infty U(s)U(t)e^{-|s-t|}\,\mathrm{d}s\,\mathrm{d}t, \\
\int_0^\infty h(\tau)\,\mathrm{d}\tau &= \sum_{k=1}^{K} \sigma_k^2 \int_0^\infty \int_0^\infty h(\tau)u_k(t)u_k(t+\tau)\,\mathrm{d}t\,\mathrm{d}\tau \\
&\quad + \frac{\gamma^2}{2}\int_0^\infty \int_0^\infty \int_0^\infty h(\tau)U(s)U(t)e^{-|s+\tau-t|}\,\mathrm{d}s\,\mathrm{d}t\,\mathrm{d}\tau,
\end{aligned}
$$

where $h(\tau) := e^{-r\tau}$. Naturally, h can be interpreted as a continuum of Lagrange multipliers, but as opposed to the discrete Lagrange multipliers, deriving h via the Euler-Lagrange equations is not feasible. Instead, inspired by numerical simulations, we guess the functional form of h. Because two solutions satisfy the Euler-Lagrange conditions, corresponding to a minimum and maximum equilibrium action, we must select the maximizer using some form of second-order condition, which is, loosely, in our setting the analogue of the classical Legendre necessary condition.

Part II: Verification

The calculus of variations determines an essentially unique candidate for the filter $\mathbf{u}$ and thus a unique candidate rating. However, few sufficient conditions are known in the calculus of variations. Most are based on the Hilbert Invariant Integral. However, in the case of (even one-dimensional) integral equations with delayed argument, the method does not apply (Sabbagh (1969)).[42] Instead, we interpret the intermediary's optimization differently, as a principal-agent model. In this auxiliary model, the agent produces signals and outputs exactly as in the original model and obtains the same payoffs. However, there is no longer a market, nor an intermediary. Instead, the agent receives transfers from a principal, who observes all outputs and signals, as does the agent. The principal's information at time t is thus $\mathcal{G}_t$, as defined in the original model. To simplify the exposition, let us focus on the confidential exclusive case. There are already two difficulties to overcome here: the action must be constant (a constraint that is difficult to formalize in the principal-agent context) and the transfer must be equal to the "market" belief.

[42] The Lagrangian can be interpreted as a bilinear quadratic form with a continuum of variables. Proving that the candidate control $\mathbf{u}$ is optimal is then equivalent to proving that the quadratic form has no saddle point. This involves a diagonalization of the quadratic form in an infinite-dimensional space, which in our case is not tractable.

The principal chooses a transfer process μ, which is interpreted as the instantaneous payment flow from the principal to the agent. As in the original model, the agent chooses an action process A (the agent's strategy) that maximizes, at all t,

$$\mathbf{E}\left[\int_{s\geq t} e^{-r(s-t)}(\mu_s - c(A_s))\,\mathrm{d}s \;\middle|\; \mathcal{G}_t\right]. \tag{17}$$

In the principal-agent formulation, the transfer process μ is not constrained to be a belief nor to have a Gaussian form.

The principal has a discount rate $\rho < r$ and seeks to maximize the *ex ante* payoff

$$\mathbf{E}\left[\int_0^\infty \rho e^{-\rho t}(c'(A_t) + \phi\mu_t(\nu_t - \mu_t))\,\mathrm{d}t\right], \tag{18}$$

where ϕ is some scalar multiplier and ν_t, defined as the mean ability of the agent under transparency, $\mathbf{E}[\theta_t \mid \mathcal{G}_t]$. The maximization is performed over all strategies A and transfer processes μ such that the action A is incentive compatible, *i.e.*, it maximizes (17).

To interpret the principal's objective, it is useful to consider the reward appearing in (18). The term $c'(A_t)$ is the agent's marginal cost, which the intermediary maximizes in the original model. If the payoff were reduced to this term, the principal might not choose a μ associated with a market belief. However, for the principal-agent and original models to be comparable, μ must be close to a market belief. The second term $\phi\mu_t(\nu_t - \mu_t)$ imposes a penalty on the principal to incite the principal to choose a μ close to a market belief. Indeed, observe that if μ_t and ν_t are jointly normal, then $\mathbf{E}\left[\mu_t(\nu_t - \mu_t)\right] = 0$ if, and only if, $\mathbf{Cov}[\mu_t, \nu_t] = \mathbf{Var}[\mu_t]$: this is the condition required for a Gaussian process μ to be a market belief, by Lemma 2.7.

If μ is a market belief process associated with a confidential information structure, then $\mathbf{Cov}[\mu_t, \nu_t] = \mathbf{Var}[\mu_t]$ and the principal's payoff is equal to

$$\mathbf{E}\left[\int_0^\infty \rho e^{-\rho t} c'(A_t)\right] = c'(A),$$

where $c'(A)$ refers to the stationary marginal cost. Thus, *the maximum payoff of the principal is never less than the marginal cost in the original model* for every ρ.

We find that there is no multiplier ϕ such that the principal maximizes his payoff by choosing a μ that is exactly a market belief. However, using the calculus of variations from Part I, we can "guess" a multiplier ϕ such that the payoff-maximizing μ approaches a market belief as $\rho \to 0$.

Note that in the original model, the intermediary must induce a constant equilibrium effort by the agent. In the principal-agent formulation, instead, the principal maximizes over all equilibrium action processes. Perhaps surprisingly, it is easier to solve this "fully dynamic" problem. Indeed, we are able to solve the principal-agent problem in closed form for every $\rho \in (0, r)$. Then, sending the principal's discount rate to zero leads to a solution that is constant in the limit, the optimal transfer tends to a market belief, and the principal's payoff becomes equal to the intermediary's objective in the original model (the agent's marginal cost). Formally, by sending ρ to 0, the maximum principal payoff converges to the conjectured maximum marginal cost from Part I. Because the principal's payoff cannot be lower than the intermediary's objective, this proves that the rating obtained in Part I is optimal.

In the public and non-exclusive cases, the methodology is similar, with a payoff specification that includes penalty terms reflecting the relevant constraints. In those cases, the principal's payoff includes additional state variables to induce the principal to choose a μ associated with public or non-exclusive market beliefs.

Note that, if we were able to properly internalize the constraint that the principal must choose transfer processes among what would correspond to market beliefs, the principal-agent formulation could, in principle, be used to obtain the necessary conditions of Part I. The difficulty is precisely that we cannot internalize these constraints, both with finite and infinite horizons, with a positive discount rate. This is why we consider a family of principal-agent problems and take limits as $\rho \to 0$. The calculus of variations then makes it possible to obtain the candidate optimal rating and the correct multipliers to be used in the principal-agent formulation.

3.4 The Incentive State as a Benchmark

To gain further insight into the role and structure of the incentive state, let us consider a special case. Suppose that signals are identical, namely, $\alpha_k = \beta_k = 1$, $\sigma_k = \sigma$ for all k. As discussed, transparency is obtained under public ratings. Let us instead consider confidential ratings. Theorem 3.1 immediately yields that, for all k,

$$u_k^c(t) = u^c(t) := \frac{1}{\sigma^2}\left[\frac{1-\sqrt{r}}{\kappa-\sqrt{r}}\sqrt{r}e^{-rt} + e^{-\kappa t}\right].$$

Hence, whether the incentive state is added or subtracted from the belief state depends on how $\sqrt{r}$ compares to 1 and κ. If $\sqrt{r}$ lies in $[1, \kappa]$, the sign of its coefficient is negative, meaning that it is subtracted. If it is outside this interval, it is added.

Plainly, which of the two impulse responses r and κ is largest depends on whether $\sqrt{r} \lesseqgtr \sqrt{\kappa} \in (1, \kappa)$, leading us to distinguish four intervals: $\sqrt{r} \in [0, 1]$, $[1, \sqrt{\kappa}]$, $[\sqrt{\kappa}, \kappa]$,

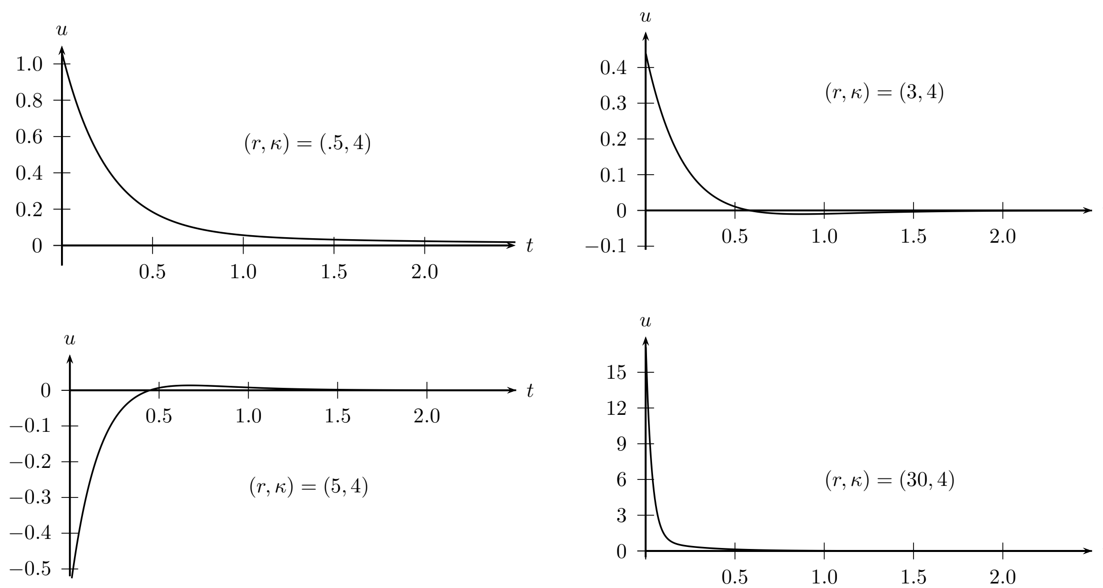

Figure 2: Rating in the case of homogeneous signals (here, $\alpha = \beta = \sigma = 1$).

and $[\kappa, \infty)$. The relative size of r vs. κ translates into how the negative sign affects the shape of $u^c(\cdot)$, as illustrated by Figure 2. If $\sqrt{r} \in [1, \sqrt{\kappa}]$, then $u(0) > 0$, but it is single-troughed and negative above some threshold t. Instead, if $\sqrt{r} \in [\sqrt{\kappa}, \kappa]$, then $u(0) < 0$ and u is single-peaked and positive above some threshold t. To see why a negative weight on the incentive term can be optimal, consider the case of a patient agent ($r < \kappa$) with output as the only signal and a rating process from the family

$$u(t) = \frac{\beta}{\sigma^2}\left(de^{-\delta t} + e^{-\kappa t}\right),$$

for some $d \in \mathbf{R}, \delta > 0$. Applying Lemma 2.12 (see (7)) yields as effort

$$c'(A) = \frac{\alpha \int_0^\infty u(t)e^{-rt}\,\mathrm{d}t}{\sqrt{\mathbf{Var}[Y_t]}}\,\mathbf{Corr}[Y_t, \theta_t] \propto \underbrace{\frac{\frac{d}{\delta+r} + \frac{1}{\kappa+r}}{\sqrt{\frac{(1+d)^2+\delta(1+d\kappa/\delta)^2}{1+\delta}}}}_{\text{Term A}}\underbrace{\mathbf{Corr}[Y_t, \theta_t]}_{\text{Term B}}. \tag{19}$$

Correlation (Term B) is maximized by transparency, setting $d = 0$: the market is never as well informed as when the intermediary reveals her own belief. Hence, to understand whether $d \gtrless 0$, we focus on the first term of (19), Term A. Its numerator is a gross (non-adjusted) measure of incentives. It is decreasing in r and linearly increasing in d: the higher the rating scale is, the greater the impact of additional effort on the rating and hence, if the market does not account for the scale, the stronger the agent's incentives. However, the market adjusts for scaling via the denominator of Term A (the standard deviation of the rating). This standard deviation is decreasing in the rate of mean-reversion (see ft. 11) and nonlinear in d. The derivative of Term A evaluated at $d = 0$ is of the same sign as

$$\frac{\kappa + r}{\delta + r} - \frac{\kappa + 1}{\delta + 1}, \tag{20}$$

the sign of which when $\delta < \kappa$ (as when $\delta = r$, its optimal value) is determined by $r \gtrless 1$. Impatience dilutes the positive impact of a higher d (the first term of (20)) on the numerator of Term A, just as mean-reversion dilutes the negative impact of a higher d (the second term) via the denominator. If impatience outweighs mean-reversion ($r > 1$), it is better to opt for a lower standard deviation and select a negative d.[43]

For $(r, \kappa) = (3, 4)$, for instance, the negative weight for $t = 1$ implies that a positive surprise at time τ negatively impacts the rating at $\tau + 1$ (see the top-right panel of Figure 2). However the rating has a positive impact until then (or, rather, until $\sim \tau + .6$). The market accounts for the fact that the rating "understates" performance; the way it is done improves its quality.

This might be a subtle point, but it is robust. While it is easiest to see in the case of identical signals, it holds for a broad range of parameters (roughly, when r is close to κ) for confidential ratings. It also occurs under public ratings, for the same reasons (indeed, it can occur under public ratings in case where it does not with confidential ratings, and vice versa). Moreover, it resonates with some practices. Murphy (2001) documents the widespread practice of past-year benchmarking as an instrument to evaluate managerial performance, commenting on its seemingly perverse incentive to underperform with an eye on the long term. Ratcheting does not explain it, as the compensation systems under study involve commitment by the firm.

3.5 Public vs. Confidential Ratings: A Closer Look

In this section, we further develop the comparison between public and confidential ratings by examining performance (effort) and informativeness (variance of the market belief). Throughout, the superscripts p and c refer to the information structure. The explicit value of the objective is given first.

[43]The case in which $r > \kappa$ can be interpreted similarly, but $c'(A)$ is not single-peaked in $d \leq 0$ in that case and the derivative at 0 is not informative.

Lemma 3.4 *The marginal cost induced is*

$$c'(A^c) = \frac{\kappa - 1}{4(\kappa + r)m_\beta} \left(2m_{\alpha\beta} + \sqrt{\Delta/r}\right),$$

given the optimal confidential rating process, and, given the optimal public process,

$$c'(A^p) = \left(1 - \left(\frac{\sqrt{r} - 1}{\sqrt{r} + 1}\right)^2\right) c'(A^c).$$

The first factor in the formula for $c'(A^p)$ quantifies the extent to which the public rating fails to match the performance of the confidential rating. Because the discount rate is the only parameter that enters this wedge, the two effort levels vary in the same way with respect to all other parameters. A higher impact of effort on signals (m_α) or noise in the type process (γ) increases effort. It is readily verified that effort is decreasing in r in the public case and that this need not be in the confidential case. In both cases, effort vanishes when $r \to \infty$ and is maximized when $r \to 0$. In the confidential case, effort then grows without bound, whereas it approaches a finite limit with a public rating. The informativeness of the rating is measured by the variance of the belief: the higher this variance is, the better informed the market.

Lemma 3.5 *The variance of the market belief is*

$$\mathbf{Var}\,\mu^c = \frac{(\kappa - 1)^2}{4m_\beta}(1 + 2m_{\alpha\beta}\sqrt{r/\Delta}),$$

given the optimal confidential rating process, and

$$\mathbf{Var}\,\mu^p = \left(1 + \left(\frac{\sqrt{r} - 1}{\sqrt{r} + 1}\right)^2\right) \mathbf{Var}\,\mu^c,$$

given the optimal public rating process.

Hence, the market is better informed given public ratings, confirming a plausible but not foregone conclusion. Here also, the wedge is a function of the discount rate alone, implying that the degrees of informativeness vary alike in all other respects.[44]

Remarkably, despite the differences in the specification between public and confidential rating systems, variance and performance differ by a constant that only

[44]Which is not to say that these comparative statics are foregone conclusions. For instance, *adding* a signal can lead to a less-informed market.

depends on the discount rate. However, with respect to the discount rate, the variation of accuracy could not be more different. As the left (right) panel of Figure 3 illustrates, variance is maximized (minimized) at an intermediate level of patience in the confidential (public) case. When ratings are confidential, an emphasis on the incentive state becomes dominant with extreme discounting. Thus, the rating becomes less accurate. Instead, given public ratings, transparency is obtained asymptotically, whether $r \to 0, \infty$. Publicness is a constraint that leaves the intermediary with little flexibility when only the long term matters ($r \approx 0$). When only the very short term matters, the incentive state decays too rapidly. As a result, under public ratings, the market backs out the belief state (the weight that the rating would have to assign to the incentive state to prevent this would transform the rating into *de facto* white noise).

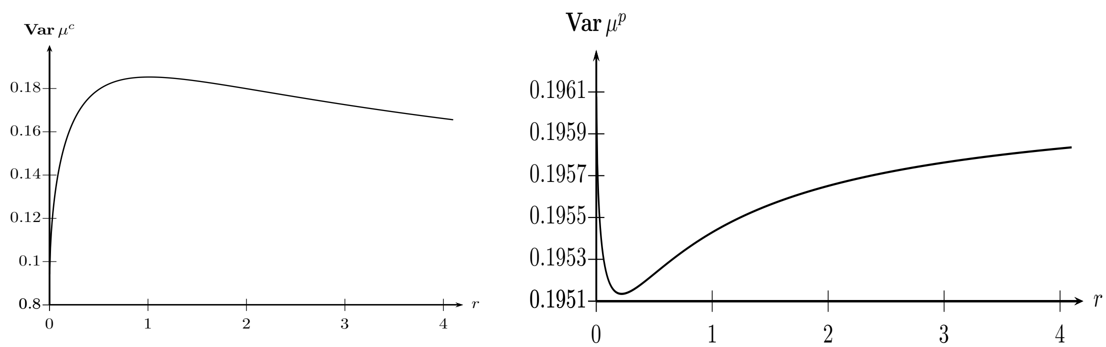

Figure 3: Confidential and public variance, as a function of r (here, $(\alpha_2, \beta_2, \sigma_2, \gamma, \sigma_1) = (3, 2, 1, 1, 2)$, $K = 2$).

This raises a natural question: is requiring ratings to be public equivalent to setting standards of accuracy? To answer this, we plot the solution (maximum marginal cost of effort) to the two problems—confidential and public ratings—*subject to* an additional constraint on the variance of the market belief. See Figure 4.[45]

Quality and effort are substitutes: transparency does not maximize effort. These substitutes are imperfect, as the effort-maximizing rating does not leave the market in the dark. Hence, there is a range of precision levels over which it conflicts with effort provision. Fixing precision, there is a maximum effort level that can be induced by the rating. (Curves are truncated at this maximum.) This maximum effort corresponds to a rating process qualitatively similar to the unconstrained one; only the weights on

[45]This plot is based on necessary conditions, but we expect that Theorems 3.1–3.2 extend.

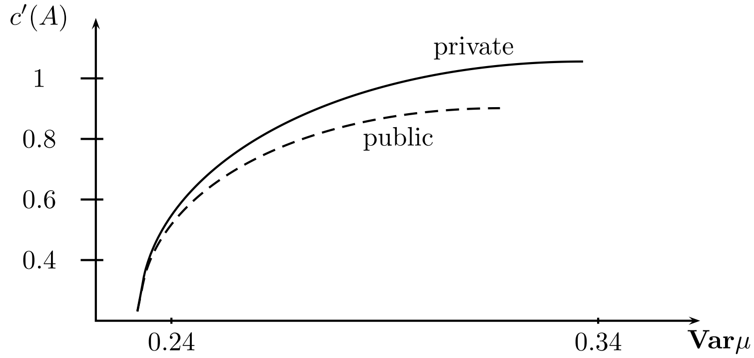

Figure 4: Marginal cost of effort as a function of maximum belief variance, public vs. confidential ratings (here, $(\beta_1, \beta_2, \alpha_1, \alpha_2, \gamma, r, \sigma_1, \sigma_2) = (3, 2, 1/3, 5, 1, 1/5, 1, 2)$).

the exponentials vary. As is clear from the figure, effort is higher in the confidential case for any given level of variance. A confidential rating system is simultaneously able to incentivize more effort and provide better information than a public system.

4 Extensions

For brevity, we focus here on two generalizations. First, we allow some signals to be non-exclusive. That is, the intermediary cannot prevent the market from observing them publicly. Second, we consider the case in which the agent's actions are multidimensional, possibly differentially affecting signals and output. All proofs are relegated to Section SA.2 of the online Supplementary Appendix.

4.1 Exclusivity

Not all information can be hidden. If the market represents long-run consumers that repeatedly interact with the agent, cumulative output is likely publicly observable. In credit ratings, solicited ratings are based on a mix of information that is widely available to market participants, as well as information that is exclusively accessible to the intermediary (see ft. 18). We refer to this distinction as *exclusive* vs. *non-exclusive* information. The intermediary does not ignore the fact that the market has direct access to this source of information. What she reveals about the exclusive signals that she can conceal also reflects the characteristics of those signals that she cannot.

Formally, all participants observe $\{S_{k,s}\}_{s\leq t,k=1,\ldots,K_0}$ in addition to the information provided by the intermediary (we consider the cases of both public and confidential structures, according to whether past information is publicly available).[46] Signals $S_{k,t}$, $k > K_0$, are only observed by the intermediary and the agent. If $K_0 = 0$, ratings are exclusive, as in Section 3. If $K_0 = K$, it is transparency, as in Section 2.4. The statements for $K_0 = 0, K$ require adjustments in the theorems given below; as they are already covered by earlier results, we rule them out.

The following proposition generalizes Proposition 2.6.

Proposition 4.1 *Let Y be a rating process. Then, Y is:*

1. *A belief for a confidential information structure with non-exclusive signals $S_1, \ldots, S_{K_0}$ if, and only if, for all t,*

$$\mathbf{E}^*[\theta_t \mid \{S_{k,s}\}_{s\leq t,k=1,\ldots,K_0}, Y_t] = Y_t.$$

2. *A belief for a public information structure with non-exclusive signals $S_1, \ldots, S_{K_0}$ if, and only if, for all t,*

$$\mathbf{E}^*[\theta_t \mid \{S_{k,s}\}_{s\leq t,k=1,\ldots,K_0}, \{Y_s\}_{s\leq t}] = Y_t.$$

The next lemma extends Lemmas 2.7 and 2.8 to account for the non-exclusive signals.

Lemma 4.2 (Confidential/Public Belief with Non-Exclusive Signals) *A rating process Y is a belief for a confidential/public information structure with non-exclusive signals $S_1, \ldots, S_{K_0}$ if, and only if, it is a belief for a confidential/public information structure and, for all $k = 1, \ldots, K_0$, all t, and all $\tau \geq 0$,*

$$\mathbf{Cov}[S_{k,t}, Y_{t+\tau}] = \mathbf{Cov}[S_{k,t}, \theta_{t+\tau}]. \tag{21}$$

Hence, the constraints that a rating process must satisfy to be a belief for a confidential or public information structure are the constraints given in Lemmas 2.7 and 2.8 respectively, and the covariance equalities (21) which capture the constraint of non-exclusivity for every signal that is publicly observed.

As in the exclusive setting of Section 3, we focus on ratings that are equal (or proportional) to beliefs, and we express the rating by its linear filter as in the Representation Lemma (Lemma 2.11). But a new choice arises: does the belief represent the interim belief based solely on the information communicated by the intermediary, to be combined with the non-exclusive signals into a posterior belief (in

[46] By our ordering convention, output is observed whenever any signal is observed.

which case, $u_k = 0$ for $k \le K_0$), or this posterior belief itself? In other words, should the rating already incorporate the information conveyed by the non-exclusive signals? This is a matter of convention. We attempt to preserve as much as possible the analogy with the solution in the exclusive case. This demands an interim approach for confidential information structures and a posterior approach for public information structures. Note that the beliefs of Lemma 4.2 refer to the posterior beliefs.

The main results of this section require the following notation. First, we introduce the rate at which a belief based solely on public signals decays, namely,

$$\hat{\kappa} := \sqrt{1 + \gamma^2 \sum_{k=1}^{K_0} \frac{\beta_k^2}{\sigma_k^2}}.$$

Second, we generalize the sums (8) to the current framework, *e.g.*,

$$m_\alpha^n := \sum_{k=1}^{K_0} \frac{\alpha_k^2}{\sigma_k^2}, \quad m_{\alpha\beta}^n := \sum_{k=1}^{K_0} \frac{\alpha_k \beta_k}{\sigma_k^2}, \quad m_\beta^n := \sum_{k=1}^{K_0} \frac{\beta_k^2}{\sigma_k^2},$$

and

$$m_\alpha^e := \sum_{k=K_0+1}^{K} \frac{\alpha_k^2}{\sigma_k^2}, \quad m_{\alpha\beta}^e := \sum_{k=K_0+1}^{K} \frac{\alpha_k \beta_k}{\sigma_k^2} \quad m_\beta^e := \sum_{k=K_0+1}^{K} \frac{\beta_k^2}{\sigma_k^2}.$$

We assume throughout that either $m_{\alpha\beta}^n \ge 0$ or $m_{\alpha\beta} \ge 0$, ensuring that positive effort can be achieved in equilibrium (by either disclosing no or all exclusive information).[47] More generally, we add superscripts n, e (for non-exclusive and exclusive) whenever convenient, with the meaning being clear from the context.

We find that Theorem 3.1 holds *verbatim*, provided we redefine Δ. Let

$$\lambda = (\kappa - 1)\left(\sqrt{r}(1+r)m_{\alpha\beta} + (\kappa^2 - r^2)\sqrt{\Delta}\right),$$

where

$$\Delta := \frac{(\kappa+1)(\hat{\kappa}+1)}{2(\kappa - \hat{\kappa})}\left[\frac{m_\alpha^e m_\beta^e}{\kappa^2 - \hat{\kappa}^2} + \frac{(1+2r+\hat{\kappa})(m_{\alpha\beta}^n)^2}{(r+\hat{\kappa})^2(\hat{\kappa}+1)} - \frac{(1+2r+\kappa)m_{\alpha\beta}^2}{(r+\kappa)^2(\kappa+1)}\right].$$

With these slightly generalized formulas, we restate Theorem 3.1.

[47] These assumptions are not necessary. The rating process defined in the theorem yields a candidate value for $c'(A)$. If it is positive, the rating system is optimal. If not, then effort is zero.

Theorem 4.3 *The optimal confidential rating process is unique and given by, for $f \leq K_0$, $u_k = 0$ and $k > K_0$,*

$$u_k(t) = \frac{\beta_k}{\sigma_k^2}\left(d_k \frac{\sqrt{r}}{\lambda} e^{-rt} + e^{-\kappa t}\right),$$

with coefficients

$$d_k := (\kappa^2 - r^2) m_\beta \frac{\alpha_k}{\beta_k} - (\kappa^2 - 1) m_{\alpha\beta}.$$

Theorem 4.4 *The optimal non-exclusive public rating process is unique and given by, for signals $k \leq K_0$,*

$$u_k^n(t) = \frac{\beta_k}{\sigma_k^2}\left(d^n e^{-\delta t} + e^{-\kappa t}\right),$$

and for signals $k > K_0$,

$$u_k^e(t) = \frac{\beta_k}{\sigma_k^2}\left(\left(c^e \frac{\beta_k}{\sigma_k^2} + d^e \frac{\alpha_k}{\beta_k}\right) e^{-\delta t} + e^{-\kappa t}\right),$$

for some constants d^n, c^e, d^e and $\delta > 0$ given in Section SA.2.5 of the online Supplementary Appendix.

The parameters d^n, c^e, d^e are elementary functions of δ, where δ is a root of a polynomial of degree 6. This polynomial is irreducible. In fact, Galois theory can be used to show that it cannot be expressed in terms of radicals. It always admits exactly two positive roots, and we indicate how to select the correct one (see Lemma SA.1 in Section SA.2.5 of the online Supplementary Appendix).

The differences in parameter values should not distract from the overarching commonalities. Most important, as in the exclusive case, the optimal process is expressed in terms of a two-state Markov process, with one state being the intermediary's belief. As before, it can be restated as a system in which the intermediary revises the rating by gradually incorporating her belief. As under exclusivity, with public ratings, the optimal rating reduces to transparency if the exclusive signals are redundant (*i.e.*, if α_k/β_k is independent of k, $k > K_0$), as is the case if there is only one such signal.

The intermediary does not need to observe the realized values of the non-exclusive signals to incentivize the agent.[48] Yet non-exclusivity affects the quality of the information available to the market. As an example, consider Figure 5, which

[48]This is not obvious from the statement of Theorem 4.4 because we chose to state the optimal non-exclusive public rating process as a posterior belief.

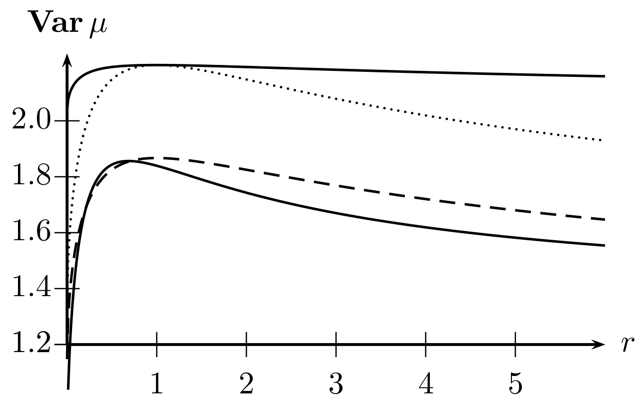

Figure 5: Belief variances (here, $K = 2$ and $(\alpha_k, \beta_k, \sigma_k, \gamma) = (1, 1, 1, 4)$, $k = 1, 2$).

describes variances under confidential ratings in a variety of cases. The market is better informed (*i.e.*, the variance of the market belief is highest) when information is non-exclusive (the higher solid line) than when it is not (the dotted line). However, this is only the case because the market can rely on the non-exclusive signal (the output) in addition to the rating. If (counterfactually) a market participant were to rely on the rating alone to derive inferences on ability (lower solid line), he would be worse off under non-exclusivity. This does not necessarily imply that the information conveyed by the rating is degraded because of the existence of another signal that the intermediary cannot hide. As is clear from Figure 5, variance could be even lower if the non-exclusive signal did not exist at all and we were considering the confidential rating process for the case of one signal only (dashed line). For nearly all discount rates, however, the presence of non-exclusive information depresses the intermediary's willingness to disclose information regarding her unshared signal—free information and the information conveyed by the rating are then strategic substitutes.[49]

4.2 Multiple Actions

Ratings are often criticized for biasing, rather than bolstering, incentives. When the agent engages in multiple tasks, a poorly designed system might distract attention from those actions that boost output and toward those that boost ratings.

[49]This is consistent with a large empirical literature in finance showing that (i) ratings do not summarize all the information that is publicly available and that (ii) the value-added of these ratings decreases in the quality of information otherwise available.

Such moral hazard takes many forms. In credit rating, for instance, both shirking and risk-shifting by the issuer are costly moral hazard activities that rating systems might encourage (see Langohr and Langohr (2009), Ch. 3). Report cards in sectors such as health care and education are widely criticized for encouraging providers to "game" the system, leading doctors to inefficient selection behavior and teachers to concentrate their effort on developing those skills measured by standardized tests.[50]

Our model can accommodate such concerns. We illustrate how in the context of confidential ratings. Suppose that there is not one but L effort levels A_ℓ, $\ell = 1, \ldots, L$, with a cost of effort that is additively separable.[51] With some abuse of notation,

$$c(A_1, \ldots, A_L) \coloneqq \sum_{\ell=1}^{L} c(A_\ell).$$

For concreteness, assume that $c(A_\ell) = cA_\ell^2$, $c > 0$, although the method applies more generally. Signals are now defined by their law

$$\mathrm{d}S_{k,t} = \left(\textstyle\sum_\ell \alpha_{k,\ell} A_{\ell,t} + \beta_k \theta_t\right) \mathrm{d}t + \sigma_k \, \mathrm{d}W_{k,t},$$

for all $k = 1, \ldots, K$, with $\sum_\ell \alpha_{1,\ell} \neq 0$. The model is otherwise unchanged. The intermediary's objective is the maximization of the expected discounted output, as in the baseline model.

This model is solved as in Section 3.2 via a change of variables. Define a fictitious model with one-dimensional effort A, cost $c(A) = cA^2$ and signals $\widetilde{S}_k$ with law

$$\mathrm{d}\widetilde{S}_{k,t} = (\alpha_k A_t + \beta_k \theta_t) \, \mathrm{d}t + \sigma_k \, \mathrm{d}W_{k,t},$$

for all $k = 1, \ldots, K$, where

$$\alpha_k \coloneqq \frac{\sum_\ell \alpha_{1,\ell} \alpha_{k,\ell}}{\sum_\ell \alpha_{1,\ell}}.$$

Lemma 4.5 *The linear filter of the optimal confidential rating process is the same in both the original model and the fictitious model.*

In terms of the optimal linear filter $\{u_k\}_k$ for the fictitious model, each effort level in the original model is then given by

$$c'(A_\ell) = \frac{\mathbf{Cov}[Y_t, \theta_t]}{\mathbf{Var}[Y_t]} \int_0^\infty e^{-rt} \left(\sum_k \alpha_{k,\ell} u_{k,t} \right) \mathrm{d}t,$$

[50] See Porter (2015) for a variety of other examples.

[51] For a discussion of the restriction implied by separability, see Holmstrom and Milgrom (1991).

for $\ell = 1, \ldots, L$. The following example shows that the optimal rating remains opaque and does not seek to deter effort in unproductive tasks. Output is only a function of effort A_1; however, the signal S_2 reflects both effort A_2 and the agent's type; namely,

$$\begin{aligned} \mathrm{d}S_{1,t} &= A_{1,t}\,\mathrm{d}t + \sigma_1\,\mathrm{d}W_{1,t}, \text{ and} \\ \mathrm{d}S_{2,t} &= (A_{2,t} + \theta_t)\,\mathrm{d}t + \sigma_2\,\mathrm{d}W_{2,t}. \end{aligned}$$

Absent any rating, if either only the first signal or both signals are observed, the unique equilibrium involves $A_\ell = 0$, $\ell = 1, 2$. Action A_1 does not affect learning about the type, and the type does not enter output. The optimal rating is given by

$$u_1(t) = \frac{\sqrt{r}}{\sigma_1} e^{-rt}, \text{ and } u_2(t) = \frac{e^{-\kappa t}}{\sigma_2^2}.$$

The signal that is irrelevant for learning is not discarded. Rather, it is exclusively assigned to the incentive term; conversely, the signal that matters for learning matters only for the learning term. This leads to positive effort on both dimensions, namely,

$$c'(A_1) = \frac{\kappa - 1}{4\sqrt{r}\sigma_1}, \quad c'(A_2) = \frac{\kappa - 1}{2(r + \kappa)\sigma_2^2},$$

and market belief variance $\frac{1}{4}(\kappa - 1)^2\sigma_2^2$. Unproductive effort in the unobservable dimension that affects learning is the price to pay for effort in the productive activity.

4.3 Performance of Standard Policies

Many real-world systems do not use two-state mixtures. Here, we illustrate how our methods also allow us to compare some standard policies that are used in practice. As mentioned in Section 2.2, exponential smoothing and moving windows are two common systems. We argue that a properly calibrated exponential smoothing rating process outperforms any moving window rating process. For simplicity, we focus on confidential exclusive ratings with only one additional signal (simply denoted S_t).

Formally, in the case of exponential smoothing, the intermediary releases signal

$$Y_t = \int_{s \le t} e^{-\delta(t-s)} \left[c\,\mathrm{d}X_s + (1 - c)\,\mathrm{d}S_s\right]$$

at time t, where $\delta > 0$ is the coefficient of smoothing and c is the relative weight placed on the output. With a moving window, the intermediary releases a signal

$$Y_t = \int_{t-T}^{t} \left[c \, \mathrm{d}X_s + (1-c) \, \mathrm{d}S_s \right],$$

where $T > 0$ is the size of the moving window. The *optimal* exponential smoothing (resp., moving window) system is defined by the choice of (c, δ) (resp., (c, T)) such that equilibrium effort is maximized. It is simple to show the following.

Lemma 4.6 *The optimal confidential exponential smoothing rating process yields higher effort than any moving window rating process.*

The proof establishes a stronger statement: for any weight c on the output, the best rating process using exponential smoothing with that weight outperforms the best moving window rating process with the same weight.

5 Concluding Comments

Our stylized model lays bare why one should not expect ratings to be Markovian and why, for instance, the same performance can have an impact on the rating that is either positive or negative according to its vintage. Richer versions might deliver more nuanced rating systems, but will not overturn these insights.

Yet, it is desirable to extend the analysis. First, in terms of technology, we assume that effort and ability are substitutes. While this follows Holmström (1999) and most of the literature on career concerns, it is limiting, as Dewatripont, Jewitt, and Tirole (1999) make clear. Building on Cisternas (2015), for instance, one might hope to relax this assumption. The absence of risk-aversion allows us to use effort as a yardstick for efficiency. Allowing for CARA preferences, for instance, would be useful to discuss the welfare implications of the informativeness of ratings.

Second, in terms of market structure, we assume a competitive market without commitment, and a single agent. When the firm that designs the rating system is the same that pays the worker, one might wish to align its ability to commit along these two dimensions. Harris and Holmstrom (1982) offer an obvious framework. Relative performance evaluation requires introducing more agents, but is also a natural extension, given the prevalence of the practice in performance appraisal.

Third, in terms of the rating process. Stationarity, in particular, is a strong assumption. It is needed (among other uses) for the Representation Lemma (Lemma 2.11), a fundamental building block. One can bypass this difficulty by simply asserting that the rating process admits (not necessarily stationary) linear filter.

References

Adomavicius, G., and A. Tuzhilin (2005): "Toward the Next Generation of Recommender Systems: A Survey of the State-of-the-Art and Possible Extensions," *IEEE Transactions on Knowledge and Data Engineering*, 177, 734–749.

Altman (1998): "The Importance and Subtlety of Credit Rating Migration," *Journal of Banking & Finance*, 22, 1231–1247.

Altman, E.I., and D.L. Kao (1992): "Rating Drift of High Yield Bonds," *Journal of Fixed Income*, 1, 15–20.

Baker G., R. Gibbons, and K.J. Murphy (1994): "Subjective Performance Measures in Optimal Incentive Contracts," *The Quarterly Journal of Economics*, 109, 1125–1156.

Bangia, A., F.X. Diebold, A. Kronimus, C. Schagen, and T. Schuermann (2002): "Ratings Migration and the Business Cycle, with Applications to Credit Portfolio Stress Testing," *Journal of Banking & Finance*, 26, 235–264.

Bar-Isaac, H., G. Caruana, and V. Cuñat (2012): "Search, Design and Market Structure," *The American Economic Review*, 102, 1140–1160.

Cantor, R., and C. Mann (2007): "Analyzing the Tradeoff between Ratings Accuracy and Stability," *Journal of Fixed Income*, 16, 60–68.

Che, Y.-K., and J. Hörner (2015): "Optimal Design for Social Learning," Cowles Discussion Paper #2000, Yale University.

Chetty, R., J. Friedman, and J. Rockoff (2014a): "Measuring the Impacts of Teachers I: Evaluating Bias in Teacher Value-Added Estimates," *The American Economic Review*, 104, 2593–2632.

—— (2014b): "Measuring the Impacts of Teachers II: Teacher Value-Added and Student Outcomes in Adulthood," *The American Economic Review*, 104, 2633–2679.

Chu, K.-C. (1972): "Team Decision Theory and Information Structures in Optimal Control Problems—Part II," *IEEE Transactions on Automatic Control*, 17, 23–28.

Cisternas, G. (2012): "Shock Persistence, Endogenous Skills and Career Concerns," Working Paper, MIT Sloan.

——(2015): "Two-Sided Learning and Moral Hazard," Working Paper, MIT Sloan.

Coffee, J.C. (1997): "Brave New World? The Impact(s) of the Internet on Modern Securities Regulation," *Business Lawyer*, 52, 1195–1233.

Dellarocas, C. (2006): "How Often Should Reputation Mechanisms Update a Trader's Reputation Profile?," *Information Systems Research*, 17, 271–285.

Dewatripont, M., I. Jewitt, and J. Tirole (1999): "The Economics of Career Concerns, Part I: Comparing Information Structures," *The Review of Economic Studies*, 66, 183–201.

Dranove, D., D. Kessler, M. McClellan, and M. Satterthwaite (2003): "Is More Information Better? The Effects of "Report Cards" on Health Care Providers," *The Journal of Political Economy*, 111, 555–588.

Ely, J.C. (2015): "Beeps," Working Paper, Northwestern University.

Ekmekci, M. (2011): "Sustainable Reputations with Rating Systems," *Journal of Economic Theory*, 151, 2–29.

Frydman, H. (2005): "Estimation in the Mixture of Markov Chains Moving with Different Speeds," *Journal of the American Statistical Association*, 100, 1046–1053.

Frydman, H., and T. Schuermann (2008): "Credit Rating Dynamics and Markov Mixture Models," *Journal of Banking & Finance*, 32, 1062–1075.

Giampieri, G., M. Davis, and M. Crowder (2005): "A Hidden Markov Model of default interaction," *Quantitative Finance*, 5, 27–34.

Glazer, J., and T.G. McGuire (2006). " Optimal quality reporting in markets for health plans," *Journal of Health Economics*, 25, 295–310.

Gonzalez F., F. Haas, R. Johannes, M. Persson, L. Toledo, R. Violi, M. Wieland, and C. Zins (2004): "Market Dynamics Associated with Credit Ratings, a Literature Review," ECB Occasional Paper Series No. 16.

Green, J., and N. Stokey (1983): "A Comparison of Tournaments and Contracts," *The Journal of Political Economy*, 91, 349–364.

Hansen, S.E. (2013): "Performance Feedback with Career Concerns," *The Journal of Law, Economics, and Organization*, 29, 1279–1316.

Harris, M., and B. Holmstrom (1982): "A Theory of Wage Dynamics," *Econometrica*, 49, 315–333.

Ho, Y.-C., and K.-C. Chu (1972): "Team Decision Theory and Information Structures in Optimal Control Problems—Part I," *IEEE Transactions on Automatic Control*, 17, 15–22.

Holmström, B. (1999): "Managerial Incentive Problems: A Dynamic Perspective," *The Review of Economic Studies*, 66, 169–182.

Holmstrom, B., and P. Milgrom (1991): "Multitask Principal-Agent Analyses: Incentive Contracts, Asset Ownership, and Job Design," *Journal of Law, Economics, and Organization*, 7, 24–52.

Jacob, B., and L. Lefgren (2005): "Principals as Agents: Subjective Performance Measurement in Education," NBER Working Paper No. 11463.

Jeulin, T., and M. Yor (1979): "Inégalités de Hardy, semimartingales, et faux-amis," in *Séminaire de Probabilités XIII 1977/78*, Lecture Notes in Mathematics, ed. by C. Dellacherie, P.-A. Meyer, and M. Weil. Heidelberg: Springer, 332–359.

Jøsang, A., R. Ismail, and C. Boyd (2007): "A Survey of Trust and Reputation Systems for Online Service Provision," *Decision Support Systems*, 43, 618–644.

Kamenskii, G.A. (2007): *Extrema of Nonlocal Functionals and Boundary Value Problems for Functional Differential Equations*. New York, U.S.A.: Nova Science.

Kremer, I., Y. Mansour, and M. Perry (2014): "Implementing the "Wisdom of the Crowd"," *The Journal of Political Economy*, 122, 998–1012.

Lando, D., and T. Skødeberg (2002): "Analyzing Ratings Transitions and Rating Drift with Continuous Observations," *Journal of Banking & Finance*, 26, 423–444.

Langohr, H., and P. Langohr (2009): *The Rating Agencies and Their Credit Ratings: What They Are, How They Work, and Why They are Relevant*. London, U.K.: Wiley Finance.

Liu, A., and A. Skrzypacz (2014): "Limited Records and Reputation Bubbles," *Journal of Economic Theory*, 146, 479–503.

Manso, G. (2013): "Feedback Effects of Credit Ratings," *Journal of Financial Economics*, 109, 535–548.

Moody's Investors Service (1991): *Global Analysis*. London, U.K.: IFR Publishing.

—— (1995): "Sovereign Supranationals Credit Opinions," September.

Morrey, C.B., Jr. (1966): *Multiple Integrals in the Calculus of Variations*. Heidelberg: Springer.

Murphy, K. (2001): "Performance Standards in Incentive Contracts," *Journal of Accounting and Economics*, 30, 245–278.

Nickell, P., W. Perraudin, and S. Varotto (2000): "Stability of Rating Transitions" *Journal of Banking & Finance*, 24, 203–227.

Pei, H.D. (2015): "Reputation with Strategic Information Disclosure," Working Paper, Massachusetts Institute of Technology.

Porter, E. (2015): "Relying on Tests in Grading the Teacher," *The New York Times*, March 24.

Portes, R. (2008): "Ratings Agency Reform," in *The First Global Financial Crisis of the 21st Century*, ed. by A. Felton and C. Reinhart. London, U.K.: VoxEU, 145–149.

Radner, R. (1961): "The Evaluation of Information in Organizations," in *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability.* Berkeley, CA: University of California Press, 491–530.

Renault, J., E. Solan, and N. Vieille (2015): "Optimal Dynamic Information Provision," `http://arxiv.org/abs/1407.5649`.

Sabbagh, L. D. (1969): "Variational Problems with Lags," *Journal of Optimization Theory and Applications*, 3, 34–51.

Stefanescu, E., R. Tunaru, and S. Turnbull (2009): "The Credit Rating Process and Estimation of Transition Probabilities: A Bayesian Approach," *Journal of Empirical Finance*, 16, 216–234.

Standard & Poor's (1994): "Sovereign Rating Criteria," Emerging Markets, October, 124–127.

A Proofs of Section 2

A.1 Proof of Lemma 2.4

1. If the cumulative payment process satisfies the zero-profit condition, then the agent who chooses effort strategy A makes (*ex ante*) payoff

$$\mathbf{E}\left[\int_0^\infty (A_t^* + \mu_t - c(A_t))\, e^{-rt}\, \mathrm{d}t\right],$$

where A^* denotes the market conjectured effort level. The agent has no impact on A^*. Thus, the agent's strategy is optimal if, and only if, it maximizes

$$\mathbf{E}\left[\int_0^\infty (\mu_t - c(A_t))\, e^{-rt}\, \mathrm{d}t\right].$$

2. If $\mathcal{F}' = \{\mathcal{F}'_t\}_{t\geq 0}$, with $\mathcal{F}'_t = \sigma(\mu_t)$, then $\mathbf{E}^*[\theta_t|\mathcal{F}'_t] = \mathbf{E}^*[\theta_t|\mu_t] = \mu_t$, hence for a given conjectured effort level A^*, the market's transfers and the agent's optimal action are the same under both information structures $\mathcal{F}$ and $\mathcal{F}'$.

A.2 Proof of Proposition 2.6

1. If Y is a belief for a confidential information structure $\mathcal{F}$, then $Y_t = \mu_t$, where, by definition, $\mu_t = \mathbf{E}^*[\theta_t \mid \mathcal{F}_t] = \mathbf{E}^*[\theta_t \mid \mu_t]$, and where the second equality follows from the law of iterated expectations. Conversely, if $Y_t = \mathbf{E}^*[\theta_t \mid Y_t]$, then Y is the belief μ for the confidential information structure induced by Y.

2. If Y is a belief for a public information structure $\mathcal{F}$, then $Y_t = \mu_t$, where, by definition, $\mu_t = \mathbf{E}^*[\theta_t \mid \mathcal{F}_t] = \mathbf{E}^*[\theta_t \mid \{\mu_s\}_{s\leq t}]$, where the second equality follows from the law of iterated expectations, using that $\mathcal{F}$ is a filtration and thus $\mathcal{F}_t$ includes all information about $\{\mu_s\}_{s\leq t}$. Conversely, if $Y_t = \mathbf{E}^*[\theta_t \mid \{Y_s\}_{s\leq t}]$, then Y is the belief μ for the public information structure that is the filtration generated by Y.

A.3 Proof of Lemma 2.7

The lemma is immediate by applying Proposition 2.6, observing that, by the projection formula for jointly normal random variables,

$$\mathbf{E}^*[\theta_t \mid Y_t] = \frac{\mathbf{Cov}[\theta_t, Y_t]}{\mathbf{Var}[Y_t]}(Y_t - \mathbf{E}^*[Y_t]).$$

A.4 Proof of Lemma 2.8

The correlation between θ_t and $\theta_{t+\tau}$ satisfies

$$\mathbf{Corr}[\theta_t, \theta_{t+\tau}] = \frac{\mathbf{Cov}[\theta_t, \theta_{t+\tau}]}{\sqrt{\mathbf{Var}[\theta_t]}\sqrt{\mathbf{Var}[\theta_{t+\tau}]}} = e^{-\tau},$$

since, as θ is a stationary Ornstein-Uhlenbeck process with reversion rate 1 and scale γ,

$$\mathbf{Cov}[\theta_t, \theta_{t+\tau}] = \frac{\gamma^2}{2}e^{-\tau}, \text{ and } \mathbf{Var}[\theta_t] = \mathbf{Var}[\theta_{t+\tau}] = \frac{\gamma^2}{2}.$$

Let μ be the market belief process induced by some public information structure $\mathcal{F}$. We have $\mathbf{E}^*[\mu_t] = \mathbf{E}^*[\theta_t] = 0$. As $\mathcal{F}$ is also a confidential information structure, μ is also a belief for a confidential information structure. Conditionally on μ_t, the random variable θ_t is then independent from every $\mu_{t-\tau}$, $\tau \geq 0$, because μ_t carries all relevant information about θ_t. Thus, $\mathbf{Cov}[\theta_t, \mu_{t-\tau} \mid \mu_t] = 0$. Let $\tau \geq 0$. The projection formulas for jointly normal random variables yield

$$\mathbf{Cov}[\theta_t, \mu_{t-\tau} \mid \mu_t] = \mathbf{Cov}[\theta_t, \mu_{t-\tau}] - \frac{\mathbf{Cov}[\theta_t, \mu_t]\,\mathbf{Cov}[\mu_{t-\tau}, \mu_t]}{\mathbf{Var}[\mu_t]}.$$

Hence,

$$\mathbf{Cov}[\mu_{t-\tau}, \mu_t] = \mathbf{Var}[\mu_t]\frac{\mathbf{Cov}[\theta_t, \mu_{t-\tau}]}{\mathbf{Cov}[\theta_t, \mu_t]} = \mathbf{Var}[\mu_{t-\tau}]\frac{\mathbf{Cov}[\theta_t, \mu_{t-\tau}]}{\mathbf{Cov}[\theta_{t-\tau}, \mu_{t-\tau}]}, \tag{22}$$

using stationarity of (μ, θ). By Lemma 2.11, there exist $u_1^\mu, \ldots, u_K^\mu$, such that μ_t can be written as

$$\mu_t = \sum_{k=1}^{K} \int_{s \leq t} u_k^\mu(t-s)[\mathrm{d}S_{k,s} - \alpha_k A_s^* \,\mathrm{d}s].$$

As $\mathbf{Cov}[\theta_t, \theta_{t-\tau}] = \gamma^2 e^{-\tau}/2$,

$$\begin{aligned}
\mathbf{Cov}[\mu_{t-\tau}, \theta_{t-\tau}] &= \frac{\gamma^2}{2} \sum_{k=1}^{K} \beta_k \int_0^\infty u_k^\mu(s) e^{-s} \,\mathrm{d}s, \text{ and} \\
\mathbf{Cov}[\mu_{t-\tau}, \theta_t] &= \frac{\gamma^2}{2} \sum_{k=1}^{K} \beta_k \int_0^\infty u_k^\mu(s) e^{-(\tau+s)} \,\mathrm{d}s = e^{-\tau} \,\mathbf{Cov}[\mu_{t-\tau}, \theta_{t-\tau}].
\end{aligned}$$

Plugging these last two expressions into (22), we have

$$\mathbf{Cov}[\mu_t, \mu_{t+\tau}] = \mathbf{Cov}[\mu_{t-\tau}, \mu_t] = \mathbf{Var}[\mu_{t-\tau}] e^{-\tau} = \mathbf{Var}[\mu_t] e^{-\tau}.$$

Now, we prove the converse. Let Y be a rating process that is a belief for a confidential information structure, and satisfies

$$\mathbf{Cov}[Y_{t+\tau}, Y_t] = \mathbf{Var}[Y_t] e^{-\tau},$$

for every $\tau \geq 0$. By Lemma 2.11, as $\mathbf{E}^*[Y_t] = 0$, there exist $u_1^Y, \ldots, u_K^Y$, such that Y_t writes as

$$Y_t = \sum_{k=1}^{K} \int_{s \leq t} u_k^Y(t-s)[\mathrm{d}S_{k,s} - \alpha_k A_s^* \,\mathrm{d}s],$$

so that, as above, we get

$$\mathbf{Cov}[Y_{t-\tau}, \theta_t] = e^{-\tau} \,\mathbf{Cov}[Y_{t-\tau}, \theta_{t-\tau}] = e^{-\tau} \,\mathbf{Cov}[Y_t, \theta_t],$$

using the stationarity of (Y, θ), and we have by assumption on Y that

$$e^{-\tau} = \frac{\mathbf{Cov}[Y_t, Y_{t-\tau}]}{\mathbf{Var}[Y_{t-\tau}]} = \frac{\mathbf{Cov}[Y_t, Y_{t-\tau}]}{\mathbf{Var}[Y_t]}.$$

Therefore,

$$\mathbf{Cov}[\theta_t, Y_{t-\tau} \mid Y_t] = \mathbf{Cov}[\theta_t, Y_{t-\tau}] - \frac{\mathbf{Cov}[\theta_t, Y_t] \,\mathbf{Cov}[Y_{t-\tau}, Y_t]}{\mathbf{Var}[Y_t]} = 0.$$

As θ and Y are jointly normal, it implies that θ_t and $Y_{t-\tau}$ are independent conditionally on Y_t for every $\tau \geq 0$, so the market belief associated with the public information structure that is the filtration generated by Y satisfies

$$\mathbf{E}^* \left[\theta_t \mid \{Y_s\}_{s \leq t}\right] = \mathbf{E}^* \left[\theta_t \mid Y_t\right] = Y_t.$$

The conclusion follows from Proposition 2.6.

A.5 Proof of Lemma 2.9 and Lemma 2.12

We prove the existence and uniqueness of the equilibrium, and give the closed-form expression of the equilibrium action. We have, using the projection formulas for jointly normal random variables, as well as the Representation Lemma (Lemma 2.11),

$$\mu_t = \mathbf{E}^*[\theta_t \mid Y_t] = \mathbf{Cov}[Y_t, \theta_t](Y_t - \mathbf{E}^*[Y_t]) = \mathbf{Cov}[Y_t, \theta_t] \sum_{k=1}^{K} \int_{s \leq t} u_k(t-s) \left[\mathrm{d}S_{k,s} - \alpha_k A_s^* \, \mathrm{d}s \right],$$

where A^* is the effort level conjectured by the market. By stationarity, $\mathbf{Cov}[Y_t, \theta_t]$ is constant. We prove that, given the (unique) cumulative payment process that satisfies the zero-profit condition, there exists a (up to measure zero sets) unique optimal effort strategy for the agent, pinned down by the first-order condition given in Lemma 2.12. This, in turn, yields existence of a unique equilibrium. Let us fix the cumulative payment process that satisfies the zero-profit condition, and suppose that the agent follows effort strategy A. The agent's time-0 (*ex post*) payoff is then

$$\int_0^\infty \left[A_t^* + \mu_t - c(A_t) \right] e^{-rt} \, \mathrm{d}t. \tag{23}$$

Maximizing the agent's *ex ante* payoff is equivalent to maximizing the *ex post* payoff (up to null events). Hence, we seek conditions on A that characterize when it is a maximizer of (23). Thus, as

$$\mathrm{d}S_{k,s} = (\alpha_k A_s + \beta_k \theta_s) \, \mathrm{d}s + \sigma_k \, \mathrm{d}Z_{k,s},$$

maximizing (23) is equivalent to maximizing

$$\mathbf{Cov}[Y_t, \theta_t] \int_0^\infty \int_0^t \sum_{k=1}^{K} \alpha_k u_k(t-s) A_s e^{-rt} \, \mathrm{d}s \, \mathrm{d}t - \int_0^\infty c(A_t) e^{-rt} \, \mathrm{d}t. \tag{24}$$

Let us re-write

$$\begin{aligned} \mathbf{Cov}[Y_t, \theta_t] \int_0^\infty \int_0^t \sum_{k=1}^{K} \alpha_k u_k(t-s) A_s e^{-rt} \, \mathrm{d}s \, \mathrm{d}t &= \mathbf{Cov}[Y_t, \theta_t] \int_0^\infty \int_s^\infty \sum_{k=1}^{K} u_k(t-s) \alpha_k A_s e^{-rt} \, \mathrm{d}t \, \mathrm{d}s \\ &= \mathbf{Cov}[Y_t, \theta_t] \int_0^\infty A_s e^{-rs} \int_0^\infty \sum_{k=1}^{K} \alpha_k u_k(\tau) e^{-r\tau} \, \mathrm{d}\tau \, \mathrm{d}s. \end{aligned}$$

Maximizing (24) is then the same as maximizing

$$\mathbf{Cov}[Y_t, \theta_t] \int_0^\infty A_s e^{-rs} \int_0^\infty \sum_{k=1}^{K} \alpha_k u_k(\tau) e^{-r\tau} \, \mathrm{d}\tau \, \mathrm{d}s - \int_0^\infty c(A_t) e^{-rt} \, \mathrm{d}t,$$

which is the same as maximizing

$$\mathbf{Cov}[Y_t, \theta_t] A_s \int_0^\infty \sum_{k=1}^{K} \alpha_k u_k(\tau) e^{-r\tau} \, \mathrm{d}\tau - c(A_s),$$

for (almost) every s. By strict convexity of the agent's cost, (24), and thus (23), is maximized if, and only if, for (almost) every t,

$$c'(A_t) = \mathbf{Cov}[Y_t, \theta_t] \int_0^\infty \sum_{k=1}^K \alpha_k u_k(\tau) e^{-r\tau} \, \mathrm{d}\tau.$$

We note that $\mathbf{Cov}[Y_t, \theta_t]$ is constant and equal to

$$\mathbf{Cov}[Y_t, \theta_t] = \frac{\gamma^2}{2} \sum_{k=1}^K \beta_k \int_0^\infty u_k(s) e^{-s} \, \mathrm{d}s.$$

Hence, (23) is maximized if, and only if, for every t (up to measure zero sets),

$$c'(A_t) = \frac{\gamma^2}{2} \left[\sum_{k=1}^K \beta_k \int_0^\infty u_k(t) e^{-t} \, \mathrm{d}t \right] \left[\sum_{k=1}^K \alpha_k \int_0^\infty u_k(t) e^{-rt} \, \mathrm{d}t \right].$$

Thus, the optimal effort strategy exists for the agent. It is unique (up to measure zero events and times). In addition, it is constant and pinned down by the last equation.

A.6 Proof of Lemma 2.11

The proof proceeds in two parts. We start with a candidate guess, whose derivation is explained in Section SA.1.1 of the online Supplementary Appendix. In the first part, we prove that the candidate obtained is integrable and square-integrable. In the second part, we show that the candidate is correct. We start with the following definitions:

$$C_1^\infty := \frac{\kappa^2 - 1}{2\kappa} \int_0^\infty \bar{f}'(j) e^{-\kappa j} \, \mathrm{d}j, \tag{25}$$

and

$$C_2^\infty := \frac{(\kappa - 1)^2}{2\kappa} \int_0^\infty \bar{f}'(j) e^{-\kappa j} \, \mathrm{d}j. \tag{26}$$

A candidate for U is

$$U(\tau) := C_1^\infty e^{\kappa\tau} + C_2^\infty e^{-\kappa\tau} - \bar{f}'(\tau) - \frac{\kappa^2 - 1}{\kappa} \int_0^\tau \sinh(\kappa(\tau - s)) \bar{f}'(s) \, \mathrm{d}s,$$

and a guess for the candidate u_k is:

$$u_k(\tau) := C_1^\infty \frac{\beta_k \gamma^2}{\sigma_k^2(\kappa^2 - 1)} e^{\kappa\tau} + C_1^\infty \frac{\beta_k \gamma^2}{\sigma_k^2(\kappa^2 - 1)} e^{-\kappa\tau} - \frac{f_k'(\tau)}{\sigma_k^2} - \frac{\beta_k \gamma^2}{\sigma_k^2 \kappa} \int_0^\tau \sinh(\kappa(\tau - s)) \bar{f}'(s) \, \mathrm{d}s, \tag{27}$$

or equivalently,

$$u_k(\tau) := \frac{\beta_k \gamma^2}{\sigma_k^2 \kappa} \left(\frac{\sinh \kappa\tau + \kappa \cosh \kappa\tau}{1 + \kappa} \int_0^\infty e^{-\kappa s} \, \mathrm{d}\bar{f}(s) - \int_0^\tau \sinh \kappa(t - s) \, \mathrm{d}\bar{f}(s) \right) - \frac{f_k'(\tau)}{\sigma_k^2}. \tag{28}$$

Proof of Integrability. We show that every u_k defined by Equation (28) is integrable and square-integrable. To do so, we have to show that

$$(\sinh \kappa t + \kappa \cosh \kappa t) \int_0^\infty e^{-\kappa s} h(s)\, \mathrm{d}s - (1+\kappa) \int_0^t \sinh \kappa(t-s) h(s)\, \mathrm{d}s \tag{29}$$

is integrable and square-integrable whenever h and h^2 are. We note that (29) is linear in h, so that it suffices to show that its positive and negative parts are integrable. Hence, without loss, we assume that $h \geq 0$. After re-arranging the terms, (29) is equal to

$$\frac{1}{2}(\kappa+1)\left(e^{\kappa t} \int_t^\infty e^{-\kappa s} h(s)\, \mathrm{d}s + e^{-\kappa t} \int_0^t e^{\kappa s} h(s)\, \mathrm{d}s\right) + \frac{1}{2}(\kappa-1) e^{-\kappa t} \int_0^\infty e^{-\kappa s} h(s)\, \mathrm{d}s. \tag{30}$$

Thus, (29) is nonnegative, and showing the integrability of (29) reduces to showing that the integral of (29) converges on $[0, +\infty)$. It is readily verified by differentiation that (29) is the derivative of

$$\frac{\cosh \kappa t + \kappa \sinh \kappa t}{\kappa} \int_0^\infty e^{-\kappa s} h(s)\, \mathrm{d}s - \frac{1+\kappa}{\kappa} \int_0^t \cosh \kappa(t-s) h(s)\, \mathrm{d}s + \frac{1+\kappa}{\kappa} \int_0^t h(s)\, \mathrm{d}s.$$

We must show that this expression converges as $t \to \infty$. Since by assumption, the last term is convergent, it suffices to show that

$$(\cosh \kappa t + \kappa \sinh \kappa t) \int_0^\infty e^{-\kappa s} h(s)\, \mathrm{d}s - (1+\kappa) \int_0^t \cosh \kappa(t-s) h(s)\, \mathrm{d}s$$

converges. Further, since

$$\cosh \kappa t + \kappa \sinh \kappa t = \frac{\kappa+1}{2} e^{\kappa t} - \frac{\kappa-1}{2} e^{-\kappa t}, \text{ and } \cosh \kappa(t-s) = \frac{e^{-\kappa(t-s)}}{2} + \frac{e^{\kappa(t-s)}}{2},$$

it suffices to show that

$$(\kappa+1) e^{\kappa t} \int_0^\infty e^{-\kappa s} h(s)\, \mathrm{d}s - (1+\kappa) \int_0^t e^{\kappa(t-s)} h(s)\, \mathrm{d}s = (\kappa+1) \int_t^\infty e^{-\kappa(s-t)} h(s)\, \mathrm{d}s$$

converges, which is immediate from the integrability of h. Thus, (29) is integrable. Next, to show that (30) is square-integrable, we show that

$$e^{\kappa t} \int_t^\infty e^{-\kappa s} h(s)\, \mathrm{d}s \tag{31}$$

is square-integrable. As square-integrable functions are closed under additivity, and h is integrable, (31) is the only term of (30) for which square-integrability is nontrivial. By the Cauchy-Schwarz inequality,

$$\left(\int_t^\infty e^{-\kappa s} h(s)\, \mathrm{d}s\right)^2 \leq \left(\int_t^\infty e^{-\kappa s} h^2(s)\, \mathrm{d}s\right) \left(\int_t^\infty e^{-\kappa s}\, \mathrm{d}s\right) = \kappa^{-1} e^{-\kappa t} \int_t^\infty e^{-\kappa s} h^2(s)\, \mathrm{d}s.$$

Thus,

$$\int_0^\tau \left(e^{\kappa t}\int_t^\infty e^{-\kappa s}h(s)\,\mathrm{d}s\right)^2 \mathrm{d}t \le \kappa^{-1}\int_0^\tau e^{\kappa t}\int_t^\infty e^{-\kappa s}h^2(s)\,\mathrm{d}s\,\mathrm{d}t$$
$$= \frac{1}{\kappa^2}\int_\tau^\infty e^{-\kappa(s-\tau)}h^2(s)\,\mathrm{d}s - \frac{1}{\kappa^2}\int_0^\infty e^{-\kappa s}h^2(s)\,\mathrm{d}s + \frac{1}{\kappa^2}\int_0^\tau h^2(t)\,\mathrm{d}t,$$

where equality follows by integration by parts. Convergence follows from square-integrability of h.

Proof that the Educated Guess is Correct. Here, we show that the candidate for $\{u_k\}_k$ defines valid coefficients for the rating process. Let u_k be defined by (28), or, equivalently, by (27). Let

$$\widetilde{Y}_t := \mathbf{E}^*[Y_t] + \sum_{k=1}^K \int_{s\le t} u_k(t-s)(\mathrm{d}S_{k,s} - \alpha_k A_s^*\,\mathrm{d}s).$$

If we have $\mathbf{Cov}[Y_t - \widetilde{Y}_t, S_{k,t-\tau}] = 0$ for all τ and k, then Y_t and $S_{k,t-\tau}$ are independent for all τ and k. As $Y_t - \widetilde{Y}_t$ is measurable with respect to the information generated by the signals $S_{k,t-\tau}$, $\tau \ge 0$, $k = 1,\dots,K$, it implies that $\mathbf{Var}[Y_t - \widetilde{Y}_t] = 0$ and thus $Y_t = \widetilde{Y}_t$. In the remainder, we show that $\mathbf{Cov}[Y_t - \widetilde{Y}_t, S_{k,t-\tau}] = 0$ for all $\tau \ge 0$ and $k = 1,\dots,K$. Let $g_k(\tau) = \mathbf{Cov}[\widetilde{Y}_t, S_{k,t-\tau}]$. We have:

$$g_k(\tau) = \sum_{i=1}^K \int_0^\infty u_i(s)\,\mathbf{Cov}[\mathrm{d}S_{i,t-s}, S_{k,t-\tau}] = \sigma_k^2\int_\tau^\infty u_k(s)\,\mathrm{d}s + \frac{\beta_k\gamma^2}{2}\int_0^\infty\int_\tau^\infty U(s)e^{-|s-j|}\,\mathrm{d}j\,\mathrm{d}s,$$

and so

$$g_k'(\tau) = -\sigma_k^2 u_k(\tau) - \frac{\beta_k\gamma^2}{2}\int_0^\infty U(s)e^{-|\tau-s|}\,\mathrm{d}s.$$

So, replacing u_k by its definition in (27),

$$\begin{aligned} g_k'(\tau) = f_k'(\tau) - C_1\frac{\beta_k\gamma^2}{\kappa^2-1}e^{\kappa\tau} - C_2\frac{\beta_k\gamma^2}{\kappa^2-1}e^{-\kappa\tau} + \frac{\beta_k\gamma^2}{\kappa}\int_0^\tau \sinh(\kappa(\tau-s))\bar{f}'(s)\,\mathrm{d}s \\ - \frac{\beta_k\gamma^2}{2}\int_0^\infty U(s)e^{-|\tau-s|}\,\mathrm{d}s. \end{aligned} \tag{32}$$

Further, multiplying (27) by β_k and summing over k, we have

$$U(\tau) = C_1^\infty e^{\kappa\tau} + C_2^\infty e^{-\kappa\tau} - \bar{f}'(\tau) - \frac{\kappa^2-1}{\kappa}\int_0^\tau \sinh(\kappa(\tau-s))\bar{f}'(s)\,\mathrm{d}s.$$

It holds that

$$\int_0^\infty U(s)e^{-|\tau-s|}\,\mathrm{d}s = \lim_{B\to\infty}\int_0^B U(s)e^{-|\tau-s|}\,\mathrm{d}s.$$

Thus,

$$\int_0^B U(s)e^{-|\tau-s|}\,\mathrm{d}s = C_1^\infty \int_0^B e^{\kappa s}e^{-|\tau-s|}\,\mathrm{d}s + C_2^\infty \int_0^B e^{-\kappa s}e^{-|\tau-s|}\,\mathrm{d}s - \int_0^B \bar{f}'(s)e^{-|\tau-s|}\,\mathrm{d}s$$
$$- \frac{\kappa^2-1}{\kappa}\int_0^B \int_0^s \sinh(\kappa(s-j))\bar{f}'(j)e^{-|\tau-s|}\,\mathrm{d}j\,\mathrm{d}s.$$

Then, for any $B > \tau$, we write

$$\int_0^B \int_0^s \sinh(\kappa(s-j))\bar{f}'(j)e^{-|\tau-s|}\,\mathrm{d}j\,\mathrm{d}s = -\frac{\kappa}{\kappa^2-1}\int_0^B \bar{f}'(j)e^{-|\tau-j|}\,\mathrm{d}j$$
$$+ \frac{e^{-B+\tau}}{\kappa^2-1}\int_0^B \bar{f}'(j)\left[\kappa\cosh(\kappa(B-j)) + \sinh(\kappa(B-j))\right]\mathrm{d}j - \frac{2}{\kappa^2-1}\int_0^\tau \sinh(\kappa(\tau-j))\bar{f}'(j)\,\mathrm{d}j.$$

Using the expressions for C_1^∞ and C_2^∞ given by (25) and (26), we get that

$$C_1^\infty\left[\frac{e^{B(\kappa-1)+\tau}}{\kappa-1} - \frac{e^{-\tau}}{\kappa+1}\right] + C_2^\infty\left[-\frac{e^{-B(\kappa+1)+\tau}}{\kappa+1} + \frac{e^{-\tau}}{\kappa-1}\right] + \frac{\kappa^2-1}{\kappa}\frac{\kappa}{\kappa^2-1}\int_0^B \bar{f}'(j)e^{-|\tau-j|}\,\mathrm{d}j$$
$$- \frac{\kappa^2-1}{\kappa}\frac{e^{-B+\tau}}{\kappa^2-1}\int_0^B \bar{f}'(j)\left[\kappa\cosh(\kappa(B-j)) + \sinh(\kappa(B-j))\right]\mathrm{d}j$$

converges to 0 as $B \to \infty$. Therefore,

$$\begin{aligned}\int_0^\infty U(s)e^{-|\tau-s|}\,\mathrm{d}s = {} & \frac{2}{\kappa}\int_0^\tau \sinh(\kappa(\tau-j))\bar{f}'(j)\,\mathrm{d}j \\ & + C_1^\infty\left[\frac{e^{\kappa\tau}}{\kappa+1} - \frac{e^{\kappa\tau}}{\kappa-1}\right] + C_2^\infty\left[\frac{e^{-\kappa\tau}}{\kappa+1} - \frac{e^{-\kappa\tau}}{\kappa-1}\right].\end{aligned} \tag{33}$$

Plugging the expression of (33) in (32) yields

$$\begin{aligned} g_k'(\tau) = {} & f_k'(\tau) - C_1^\infty\frac{\beta_k\gamma^2}{\kappa^2-1}e^{\kappa\tau} - C_2^\infty\frac{\beta_k\gamma^2}{\kappa^2-1}e^{-\kappa\tau} + \frac{\beta_k\gamma^2}{\kappa}\int_0^\tau \sinh(\kappa(\tau-s))\bar{f}'(s)\,\mathrm{d}s \\ & - \frac{\beta_k\gamma^2}{\kappa}\int_0^\tau \sinh(\kappa(\tau-j))\bar{f}'(j)\,\mathrm{d}j - \frac{\beta_k\gamma^2}{2}C_1^\infty\left[\frac{e^{\kappa\tau}}{\kappa+1} - \frac{e^{\kappa\tau}}{\kappa-1}\right] \\ & - \frac{\beta_k\gamma^2}{2}C_2^\infty\left[\frac{e^{-\kappa\tau}}{\kappa+1} - \frac{e^{-\kappa\tau}}{\kappa-1}\right] = f_k'(\tau). \end{aligned}$$

So $g_k' = f_k'$. As $f_k(0) = g_k(0) = 0$, it follows that $f = g$. Uniqueness of the coefficients (up to measure zero sets) is immediate by linearity, as different coefficients on a set of positive measure yield a different joint distributions over ratings and signals.

A.7 Proof of Theorem 2.13

Following Equation (10) of Section 2.4, we have the following linear representation of ν,

$$\nu_t = \sum_{k=1}^{K} \int_{s \le t} u_k(t-s) \left(\mathrm{d}S_{k,s} - \alpha_k A_s^* \, \mathrm{d}s \right),$$

with

$$u_k(\tau) = (\kappa - 1) \frac{\beta_k}{\sigma_k^2 m_\beta} e^{-\kappa \tau}.$$

We apply Lemma 2.12 and Equation (7), and get equilibrium action A given by

$$c'(A_t) = (\kappa - 1) \sum_{k=1}^{K} \frac{\alpha_k \beta_k}{\sigma_k^2 m_\beta} \int_0^\infty e^{-\kappa \tau} e^{-r\tau} \, \mathrm{d}\tau = (\kappa - 1) \frac{m_{\alpha\beta}}{m_\beta} \frac{1}{\kappa + r}.$$

A.8 Proof of Lemma 2.15

Let $\mathcal{F}$ be a public or confidential information structure. Let Y be a rating process proportional to the market belief, and let A be the (stationary) effort level it induces.

To show that any action in the range $[0, A]$ can be attained in the equilibrium of an alternative public/confidential information structure, we modify the rating process that achieves A to depress incentives to any desired extent. To do so, we use a source of independent noise. In addition to the K signals described in the model, we include one additional signal, indexed by $K+1$, which is entirely uninformative about both the agent's action and the agent's ability. Let us assume S_{K+1} is a two-sided standard Brownian motion. Consider the two-sided process

$$\xi_t = \int_{s \le t} e^{-(t-s)} \, \mathrm{d}S_{K+1,s}.$$

From Proposition 2.12, if Y has linear filter $\{u_k\}_k$, the equilibrium action A in both the public and confidential cases is the solution to

$$c'(A) = \frac{\mathbf{Cov}[Y_t, \theta_t]}{\mathbf{Var}[Y_t]} \sum_{k=1}^{K} \alpha_k \int_0^\infty u_k(\tau) e^{-r\tau} \, \mathrm{d}\tau.$$

Consider the alternative rating process $\widehat{Y} = (1-a)Y + a\xi$, for some $a \in [0,1]$, which is a well-defined rating process for the information generated by the $K+1$ signals. Consider the information structure generated by the rating process $\widehat{Y}$, and the induced equilibrium action, $\widehat{A}$. We have

$$\begin{aligned}
c'(\widehat{A}) &= \frac{\mathbf{Cov}[\widehat{Y}_t, \theta_t]}{\mathbf{Var}[\widehat{Y}_t]} \sum_{k=1}^{K} \alpha_k \int_0^\infty u_k(\tau) e^{-r\tau} \, \mathrm{d}\tau \\
&= \frac{(1-a)\, \mathbf{Cov}[Y_t, \theta_t]}{(1-a)^2\, \mathbf{Var}[Y_t] + a^2\, \mathbf{Var}[\xi]} \sum_{k=1}^{K} \alpha_k \int_0^\infty u_k(\tau) e^{-r\tau} \, \mathrm{d}\tau \\
&= \frac{1-a}{(1-a)^2 + a^2\, \mathbf{Var}[\xi]/\mathbf{Var}[Y_t]} c'(A).
\end{aligned}$$

By varying a over the interval $[0, 1]$, $c'(\widehat{A})$ covers the entire interval $[0, c'(A)]$, and thus $\widehat{A}$ covers the interval $[0, A]$. Besides, as Y and ξ are independent, for $\tau \geq 0$,

$$\mathbf{Cov}[\widehat{Y}_t, \widehat{Y}_{t+\tau}] = (1-a)^2\, \mathbf{Cov}[Y_t, Y_{t+\tau}] + a^2\, \mathbf{Cov}[\xi_t, \xi_{t+\tau}].$$

By Itô's isometry, we get

$$\mathbf{Cov}[\xi_t, \xi_{t+\tau}] = \int_0^\infty e^{-s} e^{-(s+\tau)} \, \mathrm{d}s = \frac{1}{2} e^{-\tau} = \mathbf{Var}[\xi_t] e^{-\tau}.$$

By Lemma 2.8, $\mathbf{Cov}[Y_t, Y_{t+\tau}] = \mathbf{Var}[Y_t] e^{-\tau}$. Thus,

$$\mathbf{Cov}[\widehat{Y}_t, \widehat{Y}_{t+\tau}] = ((1-a)^2\, \mathbf{Var}[Y_t] + a^2\, \mathbf{Var}[\xi_t]) e^{-\tau} = \mathbf{Var}[\widehat{Y}_t] e^{-\tau}.$$

Invoking Lemma 2.8 again, $\widehat{Y}$ is proportional to the market belief of a public information structure. Hence, $\widehat{A}$ also denotes the equilibrium action under that structure. Thus, under both the public and confidential information structure, any action in $[0, A]$ can be induced in equilibrium.

A.9 Proof of Lemma 2.16

We note that θ_t and μ_t are jointly normal, and as μ_t is the market belief, $\mathbf{Cov}[\theta_t, \mu_t] = \mathbf{Var}[\mu_t]$ by Lemma 2.7, so applying the projection formulas, we obtain

$$\mathbf{Var}[\theta_t \mid \mu_t] = \mathbf{Var}[\theta_t] - \frac{\mathbf{Cov}[\theta_t, \mu_t]^2}{\mathbf{Var}[\mu_t]} = \frac{\gamma^2}{2} - \mathbf{Var}[\mu_t].$$

Appendix B: Proofs of Section 3

B.1 Proof of Theorem 3.1

In this section, we prove Theorem 3.1. In Section SA.1.2 of the online Supplementary Appendix, we use calculus of variations to obtain a candidate optimal rating. Here, we focus on the verification, which is self-contained, and establishes that the candidate rating policy is optimal. To so so, we consider an auxiliary principal-agent setting. We refer to the principal-agent setting as the auxiliary setting, and to the main setting detailed in Section 2 as the original setting.

Auxiliary Setting. In the auxiliary setting, there is a principal (she) and an agent (he). Time $t \geq 0$ is continuous and the horizon infinite. The agent is as in the original model. He exerts private effort (his action), has an exogenous random ability, produces output X and generates signals $S_1 = X$, $S_2, \ldots, S_K$ over time. The various laws of motion, for the agent's ability, output, signals, are as in the original setting. The filtration $\mathcal{G}$ captures all information of the signal processes as in the original setting. The agent's information at time t continues to be $\mathcal{G}_t$, as defined in Section 2.1. The agent's strategy, which specifies his private action at every moment as a function of his information, continues to be a bounded $\mathcal{G}$-adapted process A.

However, the agent's payoff is not defined as in Section 2.1. In the auxiliary setting, the agent is not paid by a market, but by a principal. Informally, over the interval $[t, t + \mathrm{d}t)$, the principal transfers the amount $Y_t \, \mathrm{d}t$ to the agent. Here, Y is a stochastic process interpreted as a transfer rate (payments may be negative). The agent is risk-neutral; he discounts future payoffs at rate $r > 0$, and his instantaneous cost of effort is $c(\cdot)$, as in the original setting. The agent's realized discounted payoff is

$$\int_0^\infty e^{-rt} \left(Y_t - c(A_t) \right) \mathrm{d}t.$$

Given Y, the agent chooses a strategy A that maximizes his expected discounted payoff, namely,

$$A \in \operatorname*{argmax}_{\widehat{A}} \mathbf{E} \left[\int_0^\infty e^{-rt} \left(Y_t - c(\widehat{A}_t) \right) \mathrm{d}t \,\middle|\, \mathcal{G}_0 \right], \tag{34}$$

where the expectation is under the law of motion defined by strategy $\widehat{A}$. A strategy that satisfies (34) is called a *best-response* to the transfer process Y.

In the auxiliary setting, the principal combines features of both the market and the intermediary in the original setting. As the market, the principal sets the transfer to the agent, and as the intermediary, she observes all the signals the agent generates over time, *i.e.*, she knows $\mathcal{G}_t$ at time t. The principal recommends a strategy to the agent, denoted A^*—the analogue of the market conjecture in the original setting. She is risk-neutral and has discount rate $\rho \in (0, r)$. Her payoff is

$$\int_0^\infty e^{-\rho t} H_t \, \mathrm{d}t.$$

For now, there is no need to specify the instantaneous payoff process H. We specialize H below as we discuss the principal's optimization program.

A *contract* for the principal is a pair (A^*, Y). The contract is *incentive compatible* if A^* is a best-response to Y. For the most part, we focus on *stationary linear contracts.* These are contracts whose transfer processes Y are affine in the past signal increments, and are stationary: that is, there exist u_k, $k = 1, \ldots, K$, such that, up to an additive constant,

$$Y_t = \sum_{k=1}^K \int_{s \leq t} u_k(t - s) \, \mathrm{d}S_{k,s}.$$

The principal wants to maximize her own payoff over all contracts that are incentive compatible. This implies that there are two optimal control problems, one embedded into the other. First, we solve the agent's problem, and then turn to the principal's problem.

The Agent's Problem. We first state conditions of incentive compatibility. The proof follows the same arguments as in Lemma 2.12.

Lemma B.1 *Let (A, Y) be a stationary linear contract. The contract is incentive compatible if, and only if,*

$$c'(A) = \sum_{k=1}^{K} \alpha_k \int_0^\infty u(t)e^{-rt}\, \mathrm{d}t.$$

As is common in principal-agent problems, to solve the principal's problem using a dynamic programming approach, we express incentive compatibility in terms of the evolution of the agent's continuation value, or equivalently, the agent's continuation transfer. In the sequel, as in the main body of the paper, $\nu_t = \mathbf{E}[\theta_t \mid \mathcal{G}_t]$ is the agent's best current estimate about his ability.

Lemma B.2 *Let (A, Y) be a stationary linear contract. If the contract is incentive compatible, then there exists constants $C_1, \ldots, C_K$, such that the agent's continuation transfer process J defined by*

$$J_t = \mathbf{E}\left[\int_{s\geq t} e^{-r(s-t)} Y_s \,\mathrm{d}s \,\middle|\, \mathcal{G}_t\right],$$

(where the expectation is taken with respect to the law induced by strategy A) satisfies the SDE

$$\mathrm{d}J_t = (rJ_t - Y_t)\,\mathrm{d}t + \sum_{k=1}^{K}\left(\xi_\beta \frac{\gamma^2}{(1+\kappa)(1+r)}\frac{\beta_k}{\sigma_k^2} + C_k\right)[\mathrm{d}S_{k,t} - (\alpha_k A_t + \beta_k \nu_t)\,\mathrm{d}t],$$

and the two transversality conditions

$$\lim_{\tau\to+\infty} \mathbf{E}[e^{-\rho\tau} J_{t+\tau} \mid \mathcal{G}_t] = 0, \text{ and}$$
$$\lim_{\tau\to+\infty} \mathbf{E}[e^{-\rho\tau} J_{t+\tau}^2 \mid \mathcal{G}_t] = 0,$$

where $\xi_\beta := \sum_{k=1}^K \beta_k C_k$. In addition, the equilibrium action is defined by $c'(A_t) = \xi_\alpha := \sum_{k=1}^K \alpha_k C_k$.

Note that transversality is with respect to the principal's discount rate, not the agent's.

Proof. Consider a stationary linear contract (A, Y), where

$$Y_t = \sum_{k=1}^{K} \int_{s\leq t} u_k(t-s)\,[\mathrm{d}S_{k,s} - \alpha_k A_s\,\mathrm{d}s].$$

Let $J_T := \mathbf{E}\left[\int_{t\geq T} e^{-r(t-T)} Y_t\,\mathrm{d}t \,\middle|\, \mathcal{G}_T\right]$. We compute

$$\int_{t\geq T} e^{-r(t-T)} Y_t\,\mathrm{d}t = \sum_{k=1}^{K}\int_{t\geq T}\int_{s\leq T} e^{-r(t-T)} u_k(t-s)\,[\mathrm{d}S_{k,s} - A_s\,\mathrm{d}s]\,\mathrm{d}t$$
$$+ \sum_{k=1}^{K}\int_{s\geq T}\int_{t\geq s} e^{-r(t-T)} u_k(t-s)\,\mathrm{d}t\,[\mathrm{d}S_{k,s} - A_s\,\mathrm{d}s].$$

Note that, for $t \geq T$, $\mathbf{E}[\theta_t \mid \mathcal{G}_T, \theta_T] = \mathbf{E}[\theta_t \mid \theta_T] = e^{-(t-T)}\theta_T$, so using the law of iterated expectations, $\mathbf{E}[\theta_t \mid \mathcal{G}_T] = \mathbf{E}[\mathbf{E}[\theta_t \mid \mathcal{G}_T, \theta_T] \mid \mathcal{G}_T] = \mathbf{E}[e^{-(t-T)}\theta_T \mid \mathcal{G}_T] = e^{-(t-T)}\nu_T$. Hence, we can compute J_T as

$$\begin{aligned}
J_T &= \sum_{k=1}^{K} \int_{t\geq T} \int_{s\leq T} e^{-r(t-T)} u_k(t-s) \left[\mathrm{d}S_{k,s} - A_s \,\mathrm{d}s\right] \mathrm{d}t \\
&\quad + \sum_{k=1}^{K} \beta_k \int_{s\geq T} \int_{t\geq s} e^{-r(t-T)} u_k(t-s) e^{-(s-T)} \nu_T \,\mathrm{d}t \\
&= \int_{t\geq T} \int_{s\leq T} e^{-r(t-T)} u(t-s) \left[\mathrm{d}S_{k,s} - A_s \,\mathrm{d}s\right] \mathrm{d}t + \frac{\nu_T}{1+r} \sum_{k=1}^{K} \beta_k \int_{\tau\geq 0} e^{-r\tau} u_k(\tau) \,\mathrm{d}\tau.
\end{aligned}$$

Now, let us define the constants $C_1, \ldots, C_K$ as

$$C_k = \int_{\tau\geq 0} e^{-r\tau} u_k(\tau) \,\mathrm{d}\tau.$$

Then

$$\begin{aligned}
\mathrm{d}J_T &= \frac{\xi_\beta}{1+r} \,\mathrm{d}\nu_T - Y_T \,\mathrm{d}T + \sum_{k=1}^{K} C_k \left[\mathrm{d}S_{k,T} - \alpha_k A_T \,\mathrm{d}T\right] + rJ_T \,\mathrm{d}T - \frac{r}{1+r}\xi_\beta \nu_T \,\mathrm{d}T \\
&= \frac{\xi_\beta}{1+r} \,\mathrm{d}\nu_T + (rJ_T - Y_T) \,\mathrm{d}T + \sum_{k=1}^{K} C_k \left[\mathrm{d}S_{k,T} - \left(\alpha_k A_T + \frac{r}{1+r}\beta_k \nu_t\right) \mathrm{d}T\right].
\end{aligned}$$

After simplification, and using that $\mathrm{d}\nu_t = -\kappa\nu_t \,\mathrm{d}t + \frac{\gamma^2}{1+\kappa} \sum_{k=1}^{K} \frac{\beta_k}{\sigma_k^2} \left[\mathrm{d}S_{k,t} - \alpha_k A_t \,\mathrm{d}t\right]$, we get

$$\mathrm{d}J_t = (rJ_t - Y_t) \,\mathrm{d}t + \sum_{k=1}^{K} \left(\xi_\beta \frac{\gamma^2}{(1+\kappa)(1+r)} \frac{\beta_k}{\sigma_k^2} + C_k\right) \left[\mathrm{d}S_{k,t} - (\alpha_k A_t + \beta_k \nu_t) \,\mathrm{d}t\right].$$

That $c'(A_t) = \xi_\alpha$ follows from Lemma B.1. ■

Lemma B.3 *Let (A, Y) be a stationary linear contract. Suppose J and $\widehat{C}_1, \ldots, \widehat{C}_K$ are $\mathcal{G}$-adapted processes, and that J satisfies the SDE*

$$\mathrm{d}J_t = (rJ_t - Y_t) \,\mathrm{d}t + \sum_{k=1}^{K} \left(\widehat{\xi}_{\beta,t} \frac{\gamma^2}{(1+\kappa)(1+r)} \frac{\beta_k}{\sigma_k^2} + \widehat{C}_{k,t}\right) \left[\mathrm{d}S_{k,t} - (\alpha_k A_t + \beta_k \nu_t) \,\mathrm{d}t\right], \tag{35}$$

and the two transversality conditions

$$\lim_{\tau\to+\infty} \mathbf{E}[e^{-\rho\tau} J_{t+\tau} \mid \mathcal{G}_t] = 0, \text{ and} \tag{36}$$

$$\lim_{\tau\to+\infty} \mathbf{E}[e^{-\rho\tau} J_{t+\tau}^2 \mid \mathcal{G}_t] = 0, \tag{37}$$

where $\widehat{\xi}_\beta \coloneqq \sum_k \beta_k \widehat{C}_k$.

Then, J_t is the agent's continuation transfer $\mathbf{E}\left[\int_{s\geq t} e^{-r(s-t)} Y_s \,\mathrm{d}s \,\middle|\, \mathcal{G}_t\right]$, the contract is incentive compatible, and the agent's equilibrium action satisfies $c'(A_t) = \sum_k \alpha_k \widehat{C}_k$.

Proof. We fix a stationary linear contract (A, Y). Let J and $\widehat{C}_1, \dots, \widehat{C}_K$ be $\mathcal{G}$-adapted processes such that J satisfies (35), subject to (36) and (37).

Integrating J yields

$$J_t - e^{-r\tau} J_{t+\tau} =$$
$$\int_t^{t+\tau} e^{-r(s-t)} \left[Y_s - \sum_{k=1}^K \left(\widehat{\xi}_{\beta,t} \frac{\gamma^2}{(1+\kappa)(1+r)} \frac{\beta_k}{\sigma_k^2} + \widehat{C}_{k,t} \right) [\mathrm{d}S_{k,t} - (\alpha_k A_t + \beta_k \nu_t)\,\mathrm{d}t] \right],$$

and using that J is $\mathcal{G}$-adapted, together with the law of iterated expectations, we get

$$\begin{aligned}
&J_t - \mathbf{E}\left[e^{-r\tau} J_{t+\tau} \,\middle|\, \mathcal{G}_t\right] \\
&= \mathbf{E}\left[\int_t^{t+\tau} e^{-r(s-t)} Y_s \,\middle|\, \mathcal{G}_t \right] \\
&\quad + \sum_{k=1}^K \mathbf{E}\left[\int_t^{t+\tau} e^{-r(s-t)} \left(\widehat{\xi}_{\beta,t} \frac{\gamma^2}{(1+\kappa)(1+r)} \frac{\beta_k}{\sigma_k^2} + \widehat{C}_{k,t} \right) [\mathrm{d}S_{k,t} - (\alpha_k A_t + \beta_k \nu_t)\,\mathrm{d}t] \,\middle|\, \mathcal{G}_t \right] \\
&= \mathbf{E}\left[\int_t^{t+\tau} e^{-r(s-t)} Y_s \,\middle|\, \mathcal{G}_t \right].
\end{aligned}$$

Taking the limit as $\tau \to +\infty$ and applying the transversality condition (36), we get $J = V$, where V is the agent's continuation transfer, namely,

$$V_t \coloneqq \mathbf{E}\left[\int_t^{\infty} e^{-r(s-t)} Y_s \,\middle|\, \mathcal{G}_t \right].$$

As in the proof of Lemma B.2, for any stationary linear contract—incentive compatible or not—and an arbitrary strategy A of the agent, we have that

$$\mathrm{d}V_t = [rV_t - Y_t]\,\mathrm{d}t + \sum_{k=1}^K \left(\xi_\beta \frac{\gamma^2}{(1+\kappa)(1+r)} \frac{\beta_k}{\sigma_k^2} + C_k \right) [\mathrm{d}S_{k,t} - (\alpha_k A_t + \beta_k \nu_t)\,\mathrm{d}t],$$

with $C_k \coloneqq \int_{\tau\geq 0} e^{-r\tau} u_k(\tau)\,\mathrm{d}\tau$. That $J = V$ implies $\widehat{C}_k = C_k$, and thus, by Lemma B.1, the contract is incentive compatible. ■

The Principal's Problem. The problem for the principal is to choose a contract (A, Y) such that two conditions are satisfied:

1. The process Y maximizes
$$\mathbf{E}\left[\int_0^{\infty} e^{-\rho t} H_t \,\mathrm{d}t \,\middle|\, \mathcal{G}_0 \right].$$
2. The contract is incentive compatible.

In the remainder of this proof, as instantaneous payoff for the principal, we use

$$H_t := c'(A_t) - \phi Y_t(Y_t - \nu_t), \tag{38}$$

where

$$\phi := \frac{\sqrt{\Delta}}{\sqrt{r}(\kappa-1)(r+\kappa)} > 0, \tag{39}$$

and $\Delta = (r+\kappa)^2(m_\alpha m_\beta - m_{\alpha\beta}^2) + (1+r)^2 m_{\alpha\beta}^2$, as defined in Section 3.2.

Remarks on the choice of the principal's payoff: In the original setting, the intermediary seeks to maximize the agent's discounted output. In a stationary setting, it is equivalent to maximizing the agent's discounted marginal cost. The marginal cost is the first term in the right-hand side of (38). However, in the original setting, the agent's incentives are driven by the market's belief process. By Proposition 2.6, the market belief process μ satisfies

$$\mu_t = \mathbf{E}[\theta_t \mid \mu_t] = \mathbf{E}[\nu_t \mid \mu_t] = \frac{\mathbf{Cov}[\mu_t, \nu_t]}{\mathbf{Var}[\mu_t]}\mu_t,$$

using the law of iterated expectations and the projection formula for jointly normal random variables. Thus $\mathbf{Cov}[\mu_t, \nu_t] = \mathbf{Var}[\mu_t]$. To make the principal's payoff in the auxiliary setting and the intermediary's objective of the original setting comparable, we include a penalty term $\phi\mu_t(\nu_t - \mu_t)$ in the principal's payoff. Note that $\mathbf{E}\left[Y_t(\nu_t - Y_t)\right] = \mathbf{Cov}[Y_t, \nu_t] - \mathbf{Var}[Y_t]$. As a Lagrangian multiplier, the parameter ϕ captures the tradeoff between the maximization of the agent's marginal cost and the penalty term, so as to constrain the transfer to be close to a market belief. Its specific value (given in (39)) is picked using the conjectured optimal rating derived in the online Supplementary Appendix.

The principal's problem is an optimal control problem with two natural state variables: the agent's estimate of his ability, ν, and the agent's continuation transfer J. The state ν appears explicitly in the principal's payoff. Recall that ν can be expressed in closed form, namely,

$$\nu_t = \frac{\gamma^2}{1+\kappa}\sum_{k=1}^{K}\frac{\beta_k}{\sigma_k^2}\int_{s\le t} e^{-\kappa(t-s)}\left[\mathrm{d}S_{k,s} - \alpha_k A_s\,\mathrm{d}s\right].$$

Thus, for $t \ge 0$, the state variable ν is determined by its initial value,

$$\nu_0 = \frac{\gamma^2}{1+\kappa}\sum_{k=1}^{K}\frac{\beta_k}{\sigma_k^2}\int_{s\le 0} e^{\kappa s}\,\mathrm{d}S_{k,s},$$

and the equation of evolution of ν, namely,

$$\mathrm{d}\nu_t = -\kappa\nu_t\,\mathrm{d}t + \frac{\gamma^2}{1+\kappa}\sum_{k=1}^{K}\frac{\beta_k}{\sigma_k^2}\left[\mathrm{d}S_{k,s} - \alpha_k A_s\,\mathrm{d}s\right].$$

The other state J does not appear explicitly in the principal's payoff, but must be controlled to ensure that the transversality conditions are satisfied—by Lemmas B.2 and B.3, these transversality conditions are necessary and sufficient to ensure that the contract is incentive compatible.

The principal's problem can then be restated as follows: the principal seeks to find a stationary linear contract (A, Y), along with processes $\widehat{C}_k$, $k = 1, \dots, K$, so as to maximize, for all t,

$$\mathbf{E}\left[\int_t^\infty \rho e^{-\rho(s-t)}(c'(A_t) - \phi Y_t(Y_t - \nu_t))\,\mathrm{d}s \;\middle|\; \mathcal{G}_t\right]$$

subject to:

1. Incentive compatibility: $c'(A_t) = \widehat{\xi}_\alpha$, where we recall that $\widehat{\xi}_\alpha = \sum_k \alpha_k \widehat{C}_k$.
2. The evolution of the agent's belief ν, given by

$$\mathrm{d}\nu_t = -\kappa\nu_t\,\mathrm{d}t + \frac{\gamma^2}{1+\kappa}\sum_{k=1}^K \frac{\beta_k}{\sigma_k^2}\left[\mathrm{d}S_{k,s} - \alpha_k A_s\,\mathrm{d}s\right].$$

3. The evolution of the agent's continuation transfer J, given by

$$\mathrm{d}J_t = (rJ_t - Y_t)\,\mathrm{d}t + \sum_{k=1}^K \left(\widehat{\xi}_{\beta,t}\frac{\gamma^2}{(1+\kappa)(1+r)}\frac{\beta_k}{\sigma_k^2} + \widehat{C}_{k,t}\right)\left[\mathrm{d}S_{k,t} - (\alpha_k A_t + \beta_k\nu_t)\,\mathrm{d}t\right],$$

where we recall that $\widehat{\xi}_\beta = \sum_k \beta_k \widehat{C}_k$.
4. The transversality conditions, given by

$$\lim_{\tau\to+\infty} \mathbf{E}[e^{-\rho\tau}J_{t+\tau} \mid \mathcal{G}_t] = 0, \text{ and}$$
$$\lim_{\tau\to+\infty} \mathbf{E}[e^{-\rho\tau}J_{t+\tau}^2 \mid \mathcal{G}_t] = 0.$$

To solve the principal's problem, we use dynamic programming. The principal maximizes

$$\mathbf{E}\left[\int_t^\infty \rho e^{-\rho(s-t)}(\xi_{\alpha,t} - \phi Y_t(Y_t - \nu_t))\,\mathrm{d}s \;\middle|\; \mathcal{G}_t\right],$$

for every t, subject to the evolution of the state variables ν and J, and the transversality conditions on J. Without the restriction to stationary linear transfer processes, the dynamic programming problem is standard. We solve the principal's problem without imposing that restriction, and verify *ex post* that the optimal transfer in this relaxed problem is indeed stationary linear.

Assume the principal's value function V is $\mathcal{C}^2(\mathbf{R}^2)$, as a function of the two states J and ν. By standard arguments, an application of Itô's Lemma yields the Hamilton-Jacobi-Bellman (HJB) equation for V, namely

$$\begin{aligned}\rho V = \sup_{y,c_1,\dots,c_K} \rho\widehat{\xi}_\alpha - \rho\phi y(y-\nu) + (rJ - y)V_J - \nu_t V_\nu + \gamma^2\frac{\kappa-1}{\kappa+1}V_{\nu\nu} \\ + \frac{(\kappa+r)\gamma^2\widehat{\xi}_\beta}{(1+\kappa)(1+r)}V_{\nu J} + \sum_k\left(\widehat{\xi}_\beta\frac{\gamma^2}{(1+\kappa)(1+r)}\frac{\beta_k}{\sigma_k} + \sigma_k c_k\right)^2 V_{JJ},\end{aligned} \tag{40}$$

where to shorten notation we have used the subscript notation for the (partial) derivatives of V, and have abused notation by using $\widehat{\xi}_\alpha$ and $\widehat{\xi}_\beta$ to denote $\sum_k \alpha_k c_k$ and $\sum_k \beta_k c_k$, respectively.

We conjecture a quadratic value function V of the form

$$V(J,\nu) = a_0 + a_1 J + a_2\nu + a_3 J\nu + a_4 J^2 + a_5\nu^2. \tag{41}$$

Using the general form of the conjectured value function (41), we can solve for $y, c_1, \ldots, c_K$ using the first-order condition. We can then plug these variables expressed as a function of the coefficients a_i's back into (40), which allows to uniquely identify the coefficients. We obtain

$$\begin{aligned}
a_0 &= -\frac{m_{\alpha\beta}^2(\kappa-1)(1+2r+\kappa)}{4m_\beta(\kappa+r)^2(2r-\rho)\phi} + \frac{m_{\alpha\beta}(\kappa-1)}{2m_\beta(\kappa+r)} + \frac{(\kappa-1)^2\phi}{2m_\beta(2+\rho)} + \frac{m_\alpha}{4(2r-\rho)\phi},\\
a_1 &= 0, a_2 = 0, a_3 = \frac{(2r-\rho)\rho\phi}{1+r}, a_4 = -(2r-\rho)\rho\phi, \text{ and } a_5 = \frac{\rho(1-r+\rho)^2\phi}{4(1+r)^2(2+\rho)}.
\end{aligned}$$

It is readily verified that the second-order condition is equivalent to $a_5 < 0$, and so it is satisfied for all $\rho < r$. After simplification, we obtain the following expressions for y and c_k:

$$\begin{aligned}
y(J,\nu) &= (2r-\rho)J + \frac{1-r+\rho}{2(1+r)}\nu, \text{ and} \\
c_k(J,\nu) &= \frac{\alpha_k}{2\sigma_k^2(2r-\rho)\phi} - \frac{\beta_k(\kappa-1)\left(m_{\alpha\beta}(1+2r+\kappa)-(r+\kappa)(2r-\rho)\phi\right)}{2\sigma_k^2 m_\beta(r+\kappa)^2(2r-\rho)\phi}.
\end{aligned} \tag{42}$$

Thus, we obtain that the optimal processes $\widehat{C}_k$ are constant, and we get the optimal transfer at time t, Y_t, as a linear function of the state variables J_t, ν_t:

$$Y_t = \begin{bmatrix} 2r-\rho \\ \frac{1-r+\rho}{2(1+r)} \end{bmatrix} \cdot \begin{bmatrix} J_t \\ \nu_t \end{bmatrix}. \tag{43}$$

We insert the optimal control Y_t back into the equations that determine the evolution of the state variables. Doing so yields a linear two-dimensional stochastic differential equation, namely

$$\mathrm{d}\begin{bmatrix} J_t \\ \nu_t \end{bmatrix} = M\begin{bmatrix} J_t \\ \nu_t \end{bmatrix} + \sum_{k=1}^{K}\begin{bmatrix} \widehat{\xi}_{\beta,t}\frac{\gamma^2}{(1+\kappa)(1+r)}\frac{\beta_k}{\sigma_k^2} + \widehat{C}_{k,t} \\ \frac{\kappa-1}{m_\beta}\frac{\beta_k}{\sigma_k^2} \end{bmatrix}\left[\mathrm{d}S_{k,t} - \alpha_k A_t\,\mathrm{d}t\right],$$

where

$$M := \begin{bmatrix} -r+\rho & -\widehat{\xi}_\beta\frac{\kappa+r}{1+r} - \frac{1-r+\rho}{2(1+r)} \\ 0 & -\kappa \end{bmatrix}.$$

The matrix M has two eigenvalues, $-(r-\rho)$ and $-\kappa$, which are generically distinct, and negative for $\rho < r$. We can write

$$\begin{bmatrix} J_t \\ \nu_t \end{bmatrix} = \sum_{k=1}^{K}\int_{s\le t}\left(\mathbf{f}_k e^{-(r-\rho)(t-s)} + \mathbf{g}_k e^{-\kappa(t-s)}\right)\left[\mathrm{d}S_{k,t} - \alpha_k A_t\,\mathrm{d}t\right],$$

where $\mathbf{f}_k$ and $\mathbf{g}_k$ are two-dimensional vectors that can be expressed in closed form as:

$$\mathbf{f}_k := \begin{bmatrix} \frac{m_\beta(r+\kappa)(r-\kappa-\rho)}{2m_\beta(r+\kappa)(2r-\rho)(r-\kappa-\rho)\phi}\frac{\alpha_k}{\sigma_k^2} + \frac{m_{\alpha\beta}(\kappa-1)(1+\kappa+\rho)}{2m_\beta(r+\kappa)(2r-\rho)(r-\kappa-\rho)\phi}\frac{\beta_k}{\sigma_k^2} \\ 0 \end{bmatrix}, \text{ and}$$

$$\mathbf{g}_k := \begin{bmatrix} -\frac{(\kappa-1)\left(m_{\alpha\beta}(1+r)^2-(r+\kappa)(2r-\rho)(r-\kappa-\rho)\phi\right)}{2m_\beta(1+r)(r+\kappa)(2r-\rho)(r-\kappa-\rho)\phi}\frac{\beta_k}{\sigma_k^2} \\ \frac{\kappa-1}{m_\beta}\frac{\beta_k}{\sigma_k^2} \end{bmatrix}.$$

Moreover, when we plug the expressions of the state variables into (42), we get a stationary linear transfer process

$$Y_t = \sum_k \int_{s\le t} u_k(t-s)\left[\mathrm{d}S_{k,s} - \alpha_k A_s\,\mathrm{d}s\right],$$

with linear filter

$$u_k(\tau) = F_k e^{-(r-\rho)\tau} + G_k e^{-\kappa\tau},$$

where

$$F_k := \frac{m_\beta(r+\kappa)(r-\kappa-\rho)}{2m_\beta(r+\kappa)(r-\kappa-\rho)\phi}\frac{\alpha_k}{\sigma_k^2} + \frac{m_{\alpha\beta}(\kappa-1)(1+\kappa+\rho)}{2m_\beta(r+\kappa)(r-\kappa-\rho)\phi}\frac{\beta_k}{\sigma_k^2}, \text{ and}$$

$$G_k := \frac{(\kappa-1)\left(m_{\alpha\beta}(1+r)+(\kappa+r)(\kappa-r+\rho)\phi\right)}{2m_\beta(r+\kappa)(\kappa-r+\rho)\phi}\frac{\beta_k}{\sigma_k^2}.$$

The equilibrium action for the agent is stationary, and given by

$$c'(A_t) = \frac{\Delta + m_{\alpha\beta}(\kappa-1)(r+\kappa)(2r-\rho)\phi}{2m_\beta(r+\kappa)^2(2r-\rho)\phi}.$$

Thus, the contract (A,Y) just defined is an optimal stationary linear contract for the principal.

Note that, as $\rho\to 0$,

$$c'(A_t) \to \frac{\kappa-1}{4(\kappa+r)m_\beta}\left(2m_{\alpha\beta} + \sqrt{\Delta/r}\right),$$

and $\{u_k\}_k$ converges to the linear filter associated with the market belief of the conjectured optimal rating (see Section SA.1.2 of the online Supplementary Appendix).

Back to the Original Model. Now, we connect the auxiliary model and the original model, and conclude the verification. We prove by contradiction that the candidate rating obtained in Section SA.1.2 of the online Supplementary Appendix is indeed optimal. We continue to consider the auxiliary model. Let (A^*,Y^*) be the incentive-compatible contract defined by

$$c'(A^*) = \frac{\kappa-1}{4(\kappa+r)m_\beta}\left(2m_{\alpha\beta} + \sqrt{\Delta/r}\right),$$

and

$$Y_t^* = \frac{(\kappa-1)\left((\kappa-1)m_{\alpha\beta}(r+1)\sqrt{r}+\sqrt{\Delta}(\kappa-r)\right)}{2\sqrt{\Delta}m_\beta(\kappa-r)} \cdot \sum_{k=1}^{K} \int_{s\le t} u_k^c(t-s)\,[\mathrm{d}S_{k,s} - \alpha_k A_s^*\,\mathrm{d}s]\,.$$

The market belief Y_t^* is defined as the conjectured optimal rating of the original setting, while A^* is the corresponding equilibrium action. Consider an information structure $\widehat{\mathcal{F}}$, generated by some rating process, that induces a stationary action $\widehat{A}$. Let $\widehat{Y} := \mathbf{E}[\theta_t \mid \widehat{\mathcal{F}}_t]$. Note that $(\widehat{A},\widehat{Y})$ is a well-defined incentive-compatible stationary linear contract. We show that $c'(A^*) \geq c'(\widehat{A})$. Let $(A^{(\rho)}, Y^{(\rho)})$ be the optimal incentive-compatible stationary linear contract defined above, as a function of the principal's discount rate ρ. Let $V^{(\rho)}$ be the corresponding principal's expected payoff. Note that, for every confidential exclusive information structure $\mathcal{F}$ generated by a rating process, the equilibrium market belief of the original setting, $\mu_t = \mathbf{E}[\theta_t \mid \mathcal{F}_t]$, satisfies $\mathbf{Cov}[\mu_t,\nu_t] = \mathbf{Var}[\mu_t]$, and thus the principal's expected payoff for contract (A^*, Y^*) is $V^* := c'(A^*)/\rho$, while the principal's expected payoff for contract $(\widehat{A},\widehat{Y})$ is $\widehat{V} := c'(\widehat{A})/\rho$.

Then, for every $\rho \in (0,r)$, the inequalities $\rho V^{(\rho)} \geq \rho\widehat{V} = c'(\widehat{A})$ must hold. However, as $\rho \to 0$, $c'(A^{(\rho)}) \to c'(A^*)$, and the linear filter of $Y^{(\rho)}$ converges pointwise to the linear filter of Y^*. Thus, $\mathbf{Cov}[Y^{(\rho)},\nu_t] - \mathbf{Var}[Y^{(\rho)}] \to 0$, which in turn implies that $\rho V^{(\rho)} \to c'(A^*)$. Hence, $c'(A^*) \geq c'(\widehat{A})$.

B.2 Proof of Theorem 3.2

We prove Theorem 3.2. As in the proof of Theorem 3.1, we build on a conjecture derived in the online Supplementary Appendix, and verify the optimality of this candidate using an auxiliary principal-agent model. We use the same auxiliary setting described in Section B.1, with the same variables and notation, except for the principal's instantaneous payoff function H. To define the principal's payoff, we introduce an additional state variable, Λ, with initial value $\Lambda_0 = 0$, and which evolves according to

$$\mathrm{d}\Lambda_t = -r\Lambda_t\,\mathrm{d}t + Y_t\,\mathrm{d}t. \tag{44}$$

Instead of using H as in equation (38), we let

$$H_t = c'(A_t) - \phi_1 Y_t(Y_t-\nu_t) - \phi_2 Y_t\left(\frac{Y_t}{1+r} - \Lambda_t\right),$$

where

$$\phi_1 := \frac{2\sqrt{\Delta}}{(1+r)(\kappa-1)(\kappa+r)}, \text{ and } \phi_2 := \frac{\sqrt{\Delta}(r-1)}{(\kappa-1)(\kappa+r)},$$

and Δ is as defined in Section 3.2, namely, $\Delta = (r+\kappa)^2(m_\alpha m_\beta - m_{\alpha\beta}^2) + (1+r)^2 m_{\alpha\beta}^2$.

Compared to the case of the confidential exclusive setting of Section B.1, we require two penalty terms to ensure that the principal's payoff (in the auxiliary setting) and the intermediary's objective (in the original setting) are comparable. As in the confidential exclusive case, the term $\phi_1 Y_t(Y_t-\nu_t)$ can be interpreted as a Lagrangian term ensuring that the optimal transfer for the principal is close to a market belief (in the sense of the original setting). The second term, $\phi_2 Y_t\left(\frac{Y_t}{1+r} - \Lambda_t\right)$, is new. It captures the public constraint: together with the first term, it ensures that the transfer for the principal is close to a market belief derived from a *public information structure*. Indeed, recall that any public market belief μ satisfies $\mathbf{Cov}[\mu_t, \mu_{t+\tau}] = \mathbf{Var}[\mu_t]e^{-\tau}$, by Lemma 2.8. If

$$\Lambda_t = \int_0^t e^{-r(t-s)} \mu_s \,\mathrm{d}s,$$

it is immediate that Λ satisfies (44) for $Y = \mu$ and

$$\mathbf{E}\left[\mu_t \Lambda_t\right] = \int_0^t e^{-r(t-s)}\, \mathbf{Cov}\left[\mu_s, \mu_t\right] \mathrm{d}s = \frac{\mathbf{Var}\left[\mu_t\right]}{1+r}\left(1 - e^{-(1+r)t}\right).$$

Thus,

$$\mathbf{E}\left[\mu_t\left(\frac{\mu_t}{1+r} - \Lambda_t\right)\right] = \frac{\mathbf{Var}\left[\mu_t\right]}{1+r} e^{-(1+r)t}.$$

As opposed to the first penalty term, this expectation does not vanish for finite values of t, because $\Lambda_0 = 0$ (more generally, as long as Λ_0 is set independently of the contract, the above expectation cannot be zero for every market belief). However, it converges exponentially to zero as t grows, and this turns out to be sufficient for our purposes. The specific values for ϕ_1 and ϕ_2 are carefully selected using the conjectured optimal rating derived from Euler-Lagrange-type necessary conditions in Section SA.1.3 of the online Supplementary Appendix.

The principal's problem is an optimal control problem with three state variables: the agent's estimate of his ability, ν, the state associated with the public constraint, Λ, and the agent's continuation transfer, J. The state variables evolve according to

$$\begin{aligned}
\mathrm{d}\nu_t &= -\kappa\nu_t \,\mathrm{d}t + \frac{\kappa-1}{m_\beta}\sum_{k=1}^K \frac{\beta_k}{\sigma_k^2}\left[\mathrm{d}S_{k,t} - \alpha_k A_t \,\mathrm{d}t\right],\\
\mathrm{d}J_t &= (rJ_t - Y_t)\,\mathrm{d}t + \sum_{k=1}^K\left(\frac{\xi_\beta}{m_\beta}\frac{\kappa-1}{1+r}\frac{\beta_k}{\sigma_k^2} + C_k\right)\left[\mathrm{d}S_{k,t} - (\alpha_k A_t + \beta_k\nu_t)\,\mathrm{d}t\right],\\
\mathrm{d}\Lambda_t &= -r\Lambda_t \,\mathrm{d}t + Y_t \,\mathrm{d}t,
\end{aligned}$$

with $\xi_\beta \coloneqq \sum_k \beta_k C_k$ and $C_k \coloneqq \int_{\tau\geq 0} e^{-r\tau} u_k(\tau)\,\mathrm{d}\tau$. As in the confidential exclusive setting considered in Section B.1, the principal's problem can be restated as follows: the principal seeks to find a stationary linear contract (A, Y), along with processes $\widehat{C}_k$, $k = 1, \ldots, K$, which maximizes, for all t,

$$\mathbf{E}\left[\int_t^\infty \rho e^{-\rho(s-t)}\left(c'(A_s) - \phi_1 Y_s(Y_s - \nu_s) - \phi_2 Y_s\left(\frac{Y_s}{1+r} - \Lambda_s\right)\right)\mathrm{d}s \,\middle|\, \mathcal{G}_t\right]$$

subject to:

1. Incentive compatibility: $c'(A_t) = \widehat{\xi}_\alpha$, where $\widehat{\xi}_\alpha \coloneqq \sum_k \alpha_k \widehat{C}_k$.
2. The evolution of the agent's belief ν, given by

$$\mathrm{d}\nu_t = -\kappa\nu_t \,\mathrm{d}t + \frac{\kappa-1}{m_\beta}\sum_{k=1}^K \frac{\beta_k}{\sigma_k^2}\left[\mathrm{d}S_{k,s} - \alpha_k A_s \,\mathrm{d}s\right].$$

3. The evolution of the state Λ, given by $\mathrm{d}\Lambda_t = -r\Lambda_t \,\mathrm{d}t + Y_t \,\mathrm{d}t$.

4. The evolution of the agent's continuation transfer J, given by

$$\mathrm{d}J_t = (rJ_t - Y_t)\,\mathrm{d}t + \sum_{k=1}^{K}\left(\widehat{\xi}_{\beta,t}\frac{\gamma^2}{(1+\kappa)(1+r)}\frac{\beta_k}{\sigma_k^2} + \widehat{C}_{k,t}\right)\left[\mathrm{d}S_{k,t} - (\alpha_k A_t + \beta_k \nu_t)\,\mathrm{d}t\right],$$

where $\widehat{\xi}_\beta \coloneqq \sum_k \beta_k \widehat{C}_k$.

5. The transversality conditions, given by

$$\lim_{\tau\to+\infty}\mathbf{E}[e^{-\rho\tau}J_{t+\tau} \mid \mathcal{G}_t] = 0, \qquad \text{and} \qquad \lim_{\tau\to+\infty}\mathbf{E}[e^{-\rho\tau}J_{t+\tau}^2 \mid \mathcal{G}_t] = 0.$$

We use dynamic programming to solve the principal's problem. The principal maximizes

$$\mathbf{E}\left[\int_t^\infty \rho e^{-\rho(s-t)}\left(\widehat{\xi}_{\alpha,s} - \phi_1 Y_s(Y_s - \nu_s) - \phi_2 Y_s\left(\frac{Y_s}{1+r} - \Lambda_s\right)\right)\mathrm{d}s \,\middle|\, \mathcal{G}_t\right],$$

for every t, subject to the evolution of the state variables and the transversality conditions. As before, we solve the principal's problem without imposing the restriction that transfer processes be stationary linear, and verify *ex post* that the optimal transfer in this relaxed problem is stationary linear. Assume the principal's value function V is $\mathcal{C}^2(\mathbf{R}^3)$, as a function of the states J, ν and Λ. By standard arguments, Itô's Lemma yields as HJB equation

$$\begin{aligned}\rho V = \sup_{y,c_1,\ldots,c_K} & \rho\widehat{\xi}_\alpha - \rho\phi_1 y(y-\nu) - \rho\phi_2 y\left(\frac{y}{1+r} - \Lambda\right) - \nu V_\nu + (rJ - y)V_J + (-r\Lambda + y)V_\Lambda \\ & + \frac{\widehat{\xi}_\beta}{m_\beta}\frac{(\kappa-1)(\kappa+r)}{1+r}V_{\nu J} + \frac{(\kappa-1)^2}{2m_\beta}V_{\nu\nu} + \frac{1}{2}\sum_{k=1}^{K}\left(\frac{\widehat{\xi}_\beta}{m_\beta}\frac{\kappa-1}{1+r}\frac{\beta_k}{\sigma_k} + \sigma_k c_k\right)^2 V_{JJ},\end{aligned} \tag{45}$$

where, as before, to shorten notation, we have used the subscript notation for the (partial) derivatives of V, and have abused notation by using $\widehat{\xi}_\alpha$ and $\widehat{\xi}_\beta$ to denote $\sum_k \alpha_k c_k$ and $\sum_k \beta_k c_k$, respectively.

We conjecture a quadratic value function V of the form

$$V(J,\nu,\Lambda) = a_0 + a_1\nu + a_2 J + a_3\Lambda + a_4\nu J + a_5\nu\Lambda + a_6 J\Lambda + a_7\nu^2 + a_8 J^2 + a_9\Lambda^2. \tag{46}$$

We plug (46) into the dynamic programming equation (45) and solve for the optimal control variables $y, c_1, \ldots, c_K$. The equation is quadratic in $(y, c_1, \ldots, c_K)$. The second-order conditions are

$$\phi_1 + \frac{\phi_2}{1+r} > 0, \text{ and} \tag{47}$$

$$a_8 < 0. \tag{48}$$

That condition (47) is satisfied is immediate by the definition of ϕ_1 and ϕ_2. Assuming momentarily that (48) holds, the first-order conditions yield as maximizers

$$y(J,\Lambda,\nu) = \frac{(a_6-2a_8)(r+1)}{2\rho((r+1)\phi_1+\phi_2)}J + \frac{(r+1)(-a_6+2a_9+\rho\phi_2)}{2\rho((r+1)\phi_1+\phi_2)}\Lambda \tag{49}$$
$$+ \frac{(r+1)(-a_4+a_5+\rho\phi_1)}{2\rho((r+1)\phi_1+\phi_2)}\nu + \frac{(a_3-a_2)(r+1)}{2\rho((r+1)\phi_1+\phi_2)},$$
$$c_k(J,\Lambda,\nu) = \frac{(\kappa-1)(m_{\alpha\beta}\rho(\kappa+2r+1)-a_4(r+1)(\kappa+r))}{2a_8 m_\beta(\kappa+r)^2}\cdot\frac{\beta_k}{\sigma_k^2} - \frac{\rho}{2a_8}\cdot\frac{\alpha_k}{\sigma_k^2}. \tag{50}$$

Note that y is affine in the three state variables, and every c_k is constant. Define

$$\bar{\rho} = \sqrt{(\rho+2)(\rho+2r)}.$$

We then plug the optimal controls in (45) to identify the coefficients $a_0,\ldots,a_9$. Contrary to the confidential exclusive case, the system is linear-quadratic. There are two sets of coefficients that satisfy the equality (45) and the second-order conditions. However, only one set of coefficients yields a state J that satisfies the transversality condition. Keeping that set of coefficients, we get:

$$a_1 = a_2 = a_3 = 0,$$
$$a_4 = \frac{\sqrt{\Delta}\rho(2r-\rho)(\rho+r+1)\left(\rho(\rho+2)+(r-1-\rho)\bar{\rho}\right)}{(\kappa-1)(\rho+2)(r-1)r(r+1)^2(\kappa+r)},$$
$$a_5 = \frac{\sqrt{\Delta}\rho^2\left((\rho+2)(r-1-\rho)(\rho+2r)+(\rho-r+1)(\rho+r+1)\bar{\rho}\right)}{(\kappa-1)(\rho+2)(r-1)r(r+1)^2(\kappa+r)},$$
$$a_6 = \frac{\sqrt{\Delta}\rho(2r-\rho)\left(2r^2-(\rho+2)r+\rho(\bar{\rho}-\rho-1)\right)}{4(\kappa-1)r^2(\kappa+r)},$$
$$a_7 = \frac{2\sqrt{\Delta}\rho(\rho-r+1)^2(\rho+r+1)^2(\rho-\bar{\rho}+r+1)}{(\kappa-1)(\rho+2)^2(r-1)^2(r+1)^4(\kappa+r)},$$
$$a_8 = -\frac{\sqrt{\Delta}\rho(2r-\rho)\left(\rho^2+\rho-\rho r+2r(r+1)\right)}{8(\kappa-1)r^2(\kappa+r)} - \frac{\sqrt{\Delta}\rho\bar{\rho}(\rho-2r)^2}{8(\kappa-1)r^2(\kappa+r)},$$
$$a_9 = \frac{\sqrt{\Delta}\rho^3(\rho-\bar{\rho}+r+1)}{8(\kappa-1)r^2(\kappa+r)}.$$

The expression for a_0 does not impact the calculations that follow. Therefore, it is omitted. Note that, if $\rho < r$, the coefficient a_8 is negative, hence (48) is satisfied, and the maximizers are determined by the first-order condition. After inserting the coefficients $a_1,\ldots,a_9$ into (49) and (50), we obtain the optimal processes $\widehat{C}_k$, which are constant (and whose expression is lengthy and omitted), as well as the optimal transfer at time t, Y_t, as a linear function of the state variables J_t,ν_t,Λ_t:

$$Y_t = \begin{bmatrix} b_1 \\ b_2 \\ b_3 \end{bmatrix} \cdot \begin{bmatrix} J_t \\ \Lambda_t \\ \nu_t \end{bmatrix}, \tag{51}$$

with

$$b_1 := \frac{(2r-\rho)\,(\bar{\rho}-\rho+2r)}{4r}, \; b_2 := \frac{\rho\,(\rho-\bar{\rho}+2r)}{4r}, \text{ and } b_3 := \frac{\big((\rho+1)^2-r^2\big)\,(\bar{\rho}-\rho-2)}{(\rho+2)(r-1)(r+1)^2}.$$

We insert the optimal control Y_t into the equations that determine the evolution of the state variables. Doing so yields a linear three-dimensional stochastic differential equation, namely

$$\mathrm{d}\begin{bmatrix} J_t \\ \Lambda_t \\ \nu_t \end{bmatrix} = M \begin{bmatrix} J_t \\ \Lambda_t \\ \nu_t \end{bmatrix} + \sum_{k=1}^{K} \begin{bmatrix} \widehat{\xi}_{\beta,t} \dfrac{\gamma^2}{(1+\kappa)(1+r)} \dfrac{\beta_k}{\sigma_k^2} + \widehat{C}_{k,t} \\ 0 \\ \dfrac{\kappa-1}{m_\beta} \dfrac{\beta_k}{\sigma_k^2} \end{bmatrix} \left[\mathrm{d}S_{k,t} - \alpha_k A_t \,\mathrm{d}t\right],$$

where

$$M := \begin{bmatrix} r-b_1 & -b_2 & -\widehat{\xi}_\beta \dfrac{\kappa+r}{1+r} - b_3 \\ b_1 & -r+b_2 & b_3 \\ 0 & 0 & -\kappa \end{bmatrix}.$$

The matrix M has three eigenvalues,

$$\delta_f := \frac{1}{4}\Big(3\rho - \bar{\rho} - \sqrt{2}\sqrt{\rho(\rho+\bar{\rho}+1) + 2r^2 - r(\rho+2\bar{\rho}-2)} - 2r\Big),$$
$$\delta_g := \frac{1}{4}\Big(3\rho - \bar{\rho} + \sqrt{2}\sqrt{\rho(\rho+\bar{\rho}+1) + 2r^2 - r(\rho+2\bar{\rho}-2)} - 2r\Big),$$

and $\delta_h := -\kappa$. As $\rho \to 0$, it holds that $\delta_f \to -\sqrt{r}$, and $\delta_g \to -r$. Hence, if ρ is close to zero (*i.e.*, $\rho < \rho_0$, for some $\rho_0 > 0$), the eigenvalues of the matrix M are distinct and negative. We can write

$$\begin{bmatrix} J_t \\ \Lambda_t \\ \nu_t \end{bmatrix} = \sum_{k=1}^{K} \int_{s \le t} \Big(\mathbf{f}_k e^{\delta_f (t-s)} + \mathbf{g}_k e^{\delta_g (t-s)} + \mathbf{h}_k e^{\delta_h (t-s)}\Big) \left[\mathrm{d}S_{k,t} - \alpha_k A_t \,\mathrm{d}t\right],$$

where $\mathbf{f}_k, \mathbf{g}_k$ and $\mathbf{h}_k$ are three-dimensional vectors that can be expressed in closed form (the expression for $\rho > 0$ is lengthy and omitted). From (51), we get

$$Y_t = \sum_{k=1}^{K} \int_{s \le t} u_k(t-s) \left[\mathrm{d}S_{k,t} - \alpha_k A_t \,\mathrm{d}t\right],$$

with $u_k(\tau) := F_k e^{\delta_f \tau} + G_k e^{\delta_g \tau} + H_k e^{\delta_h \tau}$, for some constants $F_k, G_k, H_k, k = 1, \ldots, K$ that depend on the parameters of the model, and, in particular, on ρ. As $\rho \to 0$, we can simplify these constants.

We obtain $G_k \to 0$, and

$$F_k \to \frac{(\kappa-1)m_\beta\left(\sqrt{r}-\kappa\right)(\kappa+r)}{m_\beta\left(\sqrt{r}+1\right)\sqrt{\frac{\Delta}{r}}\left(\sqrt{r}-\kappa\right)}\frac{\alpha_k}{\sigma_k^2} + \frac{(\kappa-1)\left(\sqrt{\Delta}+(\kappa-1)m_{\alpha\beta}\left(\kappa+r-\sqrt{r}+1\right)-\sqrt{\Delta r}\right)}{m_\beta\left(\sqrt{r}+1\right)\sqrt{\frac{\Delta}{r}}\left(\sqrt{r}-\kappa\right)}\frac{\beta_k}{\sigma_k^2},$$

$$H_k \to -\frac{(\kappa-1)\left((\kappa-1)m_{\alpha\beta}(r+1)\sqrt{r}+\sqrt{\Delta}(\kappa-r)\right)}{\sqrt{\Delta}m_\beta\left(\sqrt{r}+1\right)\left(\sqrt{r}-\kappa\right)}\frac{\beta_k}{\sigma_k^2}.$$

Also, as $\rho \to 0$,

$$\widehat{C}_k \to \frac{m_\beta(\kappa+r)^2}{\sqrt{\Delta}m_\beta\left(\sqrt{r}+1\right)^2(\kappa+r)}\frac{\alpha_k}{\sigma_k^2} + \frac{(\kappa-1)\left(2\sqrt{\Delta r}-(\kappa-1)m_{\alpha\beta}(\kappa+2r+1)\right)}{\sqrt{\Delta}m_\beta\left(\sqrt{r}+1\right)^2(\kappa+r)}\frac{\beta_k}{\sigma_k^2}.$$

Thus,

$$\widehat{\xi}_\alpha(=c'(A)) \to \frac{(\kappa-1)\left(m_\alpha m_\beta(\kappa+r)^2-(\kappa-1)m_{\alpha\beta}^2(\kappa+2r+1)+2m_{\alpha\beta}\sqrt{\Delta r}\right)}{\sqrt{\Delta}m_\beta\left(\sqrt{r}+1\right)^2(\kappa+r)},$$

and so, after simplification,

$$c'(A_t) \to \frac{(\kappa-1)\left(1-\left(\frac{\sqrt{r}-1}{\sqrt{r}+1}\right)^2\right)\left(2m_{\alpha\beta}+\sqrt{\frac{\Delta}{r}}\right)}{4m_\beta(\kappa+r)}.$$

Back to the Original Model. We conclude the verification and connect the auxiliary model and the original model. The procedure is analogous to the confidential case. See the last part of Section B.1. Let (A^*, Y^*) be the incentive-compatible contract defined by

$$c'(A_t^*) = \frac{(\kappa-1)}{4m_\beta(\kappa+r)}\left(1-\left(\frac{\sqrt{r}-1}{\sqrt{r}+1}\right)^2\right)\left(2m_{\alpha\beta}+\sqrt{\frac{\Delta}{r}}\right),$$

and

$$Y_t^* = -\frac{(\kappa-1)\left((\kappa-1)m_{\alpha\beta}(r+1)\sqrt{r}+\sqrt{\Delta}(\kappa-r)\right)}{\sqrt{\Delta}m_\beta\left(\sqrt{r}+1\right)\left(\sqrt{r}-\kappa\right)} \cdot \sum_{k=1}^{K}\int_{s\le t} u_k^p(t-s)\left[\mathrm{d}S_{k,s}-\alpha_k A_s^*\,\mathrm{d}s\right].$$

Here, Y_t^* is the market belief of the conjectured optimal rating of the original setting, and A_t^* is the conjectured optimal action. Let $\widehat{\mathcal{F}}$ be a public information structure, generated by some rating process, which induces a constant action process $\widehat{A}$. Let $\widehat{Y} := \mathbf{E}[\theta_t \mid \widehat{\mathcal{F}}_t]$. Observe that $(\widehat{A}, \widehat{Y})$ is an incentive-compatible stationary linear contract. We show that $c'(A^*) \ge c'(\widehat{A})$. For $\rho < \rho_0$, let $(A^{(\rho)}, Y^{(\rho)})$ be the optimal incentive-compatible stationary linear contract defined above.

Let V^* be the principal's expected payoff under contract (A^*, Y^*), $\widehat{V}$ her expected payoff under $(\widehat{A}, \widehat{Y})$, and $V^{(\rho)}$ her expected payoff $(A^{(\rho)}, Y^{(\rho)})$. For every public exclusive information structure $\mathcal{F}$ generated by some rating process, the equilibrium market belief of the original setting, $\mu_t = \mathbf{E}[\theta_t \mid \mathcal{F}_t]$, satisfies $\mathbf{Cov}[\mu_t, \nu_t] = \mathbf{Var}[\mu_t]$, and $\mathbf{Cov}[\mu_t, \mu_{t+\tau}] = \mathbf{Var}[\mu_t]e^{-\tau}$ for all $\tau > 0$, by Lemma 2.8. Thus, under (μ, A), with A the equilibrium action, the state variable Λ is expressed as $\Lambda_t = \int_0^t e^{-r(t-s)}\mu_s \,\mathrm{d}s$, and the principal's payoff is

$$\int_0^\infty e^{-\rho t}\left(c'(A) - \phi_1\mu_t(\mu_t - \nu_t) - \phi_2\mu_t\left(\frac{\mu_t}{1+r} - \Lambda_t\right)\right)\mathrm{d}t = \frac{c'(A_t)}{\rho} - \phi_2\frac{\mathbf{Var}\left[\mu_t\right]}{(1+r)(1+r+\rho)}.$$

Hence, as $\rho \to 0$, $\rho V^* \to c'(A^*)$, and $\rho\widehat{V} \to c'(\widehat{A})$. For every ρ small enough, $V^{(\rho)} \geq \widehat{V}$ must hold, because $(A^{(\rho)}, Y^{(\rho)})$ is optimal. However, as $\rho \to 0$, $c'(A^{(\rho)}) \to c'(A^*)$, and the linear filter of $Y^{(\rho)}$ converges pointwise to the linear filter of Y^*. In particular, $\mathbf{Cov}[Y^{(\rho)}, \nu_t] - \mathbf{Var}[Y^{(\rho)}] \to 0$, and, for every $\tau > 0$, $\mathbf{Cov}[Y_t^{(\rho)}, Y_{t+\tau}^{(\rho)}] - \mathbf{Var}[Y_t^{(\rho)}]e^{-\tau} \to 0$. Together, these two limits imply that, as $\rho \to 0$, $\rho V^{(\rho)} - \rho V^* \to 0$. Thus, $\rho V^{(\rho)} \to c'(A^*)$, implying that $c'(A^*) \geq c'(\widehat{A})$.

B.3 Proof of Lemma 3.4 and Lemma 3.5

Fix a confidential or a public information structure $\mathcal{F}$. Given a rating process Y that is proportional to a market belief induced by $\mathcal{F}$ and with linear filter $\{u_k\}_k$, the equilibrium marginal cost is given by

$$c'(A) = \frac{\mathbf{Cov}[Y_t, \theta_t]}{\mathbf{Var}[Y_t]}\sum_{k=1}^{K}\alpha_k\int_0^\infty u_k(\tau)e^{-r\tau}\,\mathrm{d}\tau,$$

and the equilibrium market belief induced by the information structure is

$$\mu_t = \mathbf{E}[\theta_t \mid \mathcal{F}_t] = \mathbf{E}[\theta_t \mid Y_t] = \frac{\mathbf{Cov}[Y_t, \theta_t]}{\mathbf{Var}[Y_t]}Y_t.$$

This follows from the projection formulas for jointly normal random variables, with

$$\mathbf{Cov}[Y_t, \theta_t] = \frac{\gamma^2}{2}\sum_{k=1}^{K}\beta_k\int_0^\infty u_k(\tau)e^{-\tau}\,\mathrm{d}\tau,$$

and

$$\mathbf{Var}[Y_t] = \sum_{k=1}^{K}\sigma_k^2\int_0^\infty u_k(\tau)^2\,\mathrm{d}\tau + \sum_{k=1}^{K}\sum_{k'=1}^{K}\int_0^\infty\int_0^\infty u_k(s)u_{k'}(s')e^{-|s-s'|}\,\mathrm{d}s\,\mathrm{d}s'.$$

Thus,

$$\mathbf{Var}[\mu_t] = \frac{\mathbf{Cov}[Y_t, \theta_t]^2}{\mathbf{Var}[Y_t]}.$$

The expressions $c'(A^c)$ and $c'(A^p)$ given in the statement of Lemma 3.4, as well as the expressions $\mathbf{Var}[\mu^c]$ and $\mathbf{Var}[\mu^p]$ given in the statement of Lemma 3.5, follow by plugging in the expressions of the linear filters for the optimal rating as described in Theorem 3.1 and Theorem 3.2. The calculations are lengthy and omitted, a detailed proof being available upon request.