Motivational Ratings*

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

May 31, 2017

Abstract

Performance evaluation ("rating") systems not only provide information to users but also motivate the rated worker. This paper solves for the optimal (effort-maximizing) rating within the standard career concerns framework. We prove that this rating is a linear function of past observations. The rating, however, is not a Markov process, but rather the sum of two Markov processes. We show how it combines information of different types and vintages. An increase in effort may adversely affect some (but not all) future ratings.

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 worker. These two goals are not necessarily aligned. Reputational incentives require the market information to be sensitive to the worker's

*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. A special thanks to Sara Bertin, former analyst at Moody's, and Laurent Clavel, head of economics at AXA IM, for insightful discussions and comments.

†Yale University, 30 Hillhouse Ave., New Haven, CT 06520, USA, and Toulouse School of Economics, 21 Allée de Brienne, 31015 Toulouse, France, joh.horner@gmail.com.

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

effort, but excessive information depresses career concerns and distorts the worker's choices.[1] The purpose of this paper is to examine how the information structure (the "rating") should be designed to motivate effort. How should different sources of information be combined? At what rate should past observations be discounted?

As we show, the optimal information structure does not involve a partition of signals of a given type or vintage. Nor is it a "noised-up" version of all information that is available. Rather, the effort-maximizing rating is a linear function of the pieces of information available to the rater. This linear function is not invertible: although it does not add noise, the optimal rating system confounds different signals.

What is the structure of this linear map? Although the environment is Markovian in the rater's own belief and the public information, the rating is not simply a function of these two variables. It is not a function of the latest rating and the rater's private signal either.[2] The time series of optimal ratings fails to satisfy the Markov property.[3]

Thus, the optimal rating is not a Markov process. As we show, it is the sum of two Markov processes (what is known as a two-state Markov mixture model). As is intuitive, the rater's own private belief matters. This is the first process. The second reflects the worker's preferences. We call it the *incentive state*. Preferences determine the impulse response of the incentive state via the worker's impatience. Past observations are discounted in the rating at a rate equal to the worker's discount rate. The optimal rating balances the rater's information, as summarized by the rater's belief, against some short-termism that is in proportion to the worker's impatience.[4]

These results hold irrespective of whether past ratings can be hidden from the market (confidential vs. public ratings) and of whether the market has access to

[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 they encourage—but gaming is also widely documented (see Jacob and Lefgren (2005), among many others).

[2] In the literature on subjective performance evaluation, this is referred to as the spillover effect, whereby the worker's current-period performance is evaluated in part based on his past performance (*e.g.*, Bol and Smith, (2011)). 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]Such failures are widely documented for credit ratings; see Section 4.

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

additional non-proprietary information (exclusive vs. non-exclusive ratings). However, these distinctions matter for the particulars of the rating. For instance, if past ratings are observable, then concealing 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 other sources of information are freely available. Instead, in the confidential case, the free information and that revealed by the rating can be substitutes. Moreover, the two-state Markov mixture rating is canonical, in the sense that any effort that can be induced given some information structure can be induced by such a rating (possibly with some added noise). Hence, such ratings are without loss, independent of the rater's objective, or her ability to commit.

Surprisingly, perhaps, we show that the rating system can count past performance against the worker's rating ("past-year benchmarking"). 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, with the market adjusting for its variance. However, when the worker's ability is not too persistent (low mean-reversion), the variance of the rating is naturally high. By assigning positive values to recent signals and negative values to older signals, 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 worker is impatient. This feature of performance evaluation is well-documented and usually referred to as leniency bias, whereby good reviews can be given for poor performance (*e.g.*, Holzbach (1978)).

Our analysis builds on the seminal model of Holmström (1999).[5] A worker exerts effort that is hidden to the market, which pays him a competitive wage. This wage is based on the market's expectation of the worker'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 rater. The rater potentially has many sources of information about the worker 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

[5]Differences between our model and that of Holmström (1999) include the continuous-time setting, mean-reversion in the type process, and a multidimensional signal structure. See Cisternas (2017a) for a specification that is similar to ours in the first two respects.

maximize the worker's incentive to exert effort.[6,7]

Normally distributed beliefs and linear payoffs enable the analysis to focus on deterministic equilibria (Markov perfect equilibria being a special case), in which actions depend only on calendar time. As we prove (Theorem 3.5), linear ratings (*i.e.*, weighted integrals of all past observations using weight functions that the rater freely chooses) are optimal: no rating delivering deterministic effort can yield higher output. Furthermore, we establish an equivalence among such ratings, Gaussian ratings, and ratings such that the market information is of deterministic quality. Hence, such ratings are without loss as long as one wants to maximize output or guarantee that the contingencies of the worker's performance do not adversely affect the quality of the information available to the market. We focus primarily on the case of stationary ratings, but we also solve for the non-stationary (finite-horizon) optimal rating and show how it converges to the stationary rating.

An example illustrating some of the ideas is presented in Section 2. In Sections 3 and 4, we begin with the case of a single signal: output. This allows us to convey many of the insights. In Section 5, we extend the main characterization to the case of multiple signals and study several extensions. First, we allow for ratings that are public—that is, the market observes not only the latest rating, but also past ratings. Then, we relax the rating's exclusivity. That is, the market has access to independent public information. We show how the optimal rating reflects the content of this free information. Third, we discuss how our results extend to the case of multiple actions.[8] 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

[6]We also examine the trade-off between the level of effort and the precision of the market's information. These two objectives feature prominently in economic analyses of ratings conducted by both practitioners and theorists. 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 may provide information to the worker himself, informing rather than simply motivating his effort choice; see Hansen (2013). Moreover, whenever evaluating performance requires input from the users, ratings must account for users' 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 conflicts of interest and their inability to commit, which motivate a broad literature.

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

encourage productive effort. The reader interested in the more formal aspects of the analysis is referred to the appendix, which contains not only additional technical ingredients, and rigorous statements, but also an overview of the main proofs. The details of these proofs can be found in the Online Appendix.

Related Literature. Foremost, our paper builds on Holmström (1999). (See also Dewatripont, Jewitt, and Tirole (1999).) His model elegantly illustrates how imperfect monitoring cultivates incentives when information is incomplete. His analysis prompts the question raised and answered in our model: what type of feedback stimulates effort? Our interest in multifaceted information is similar to that of 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 in 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 (2016) 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 (2016), 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. In an interesting recent paper, Rodina (2016) shows how transparency is the optimal information structure under a wide class of technologies (output need not be a linear Gaussian function of effort and type). To the extent that his analysis is one-shot, it does not account for the fact that transparency also leads to reduced uncertainty over ability in the future, which makes it more difficult to sustain subsequent incentives. Therefore, one reading of his results is that this channel is important for the role of obfuscation.

Finally, in a paper that involves no incomplete information, Fong and Li (2017) show how in the context of relational contracts, it can be sensible to use mechanisms that spread the reward for the worker's good performance over time and display leniency bias by showing how the optimal contract within a particular family displays these features.

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, which are 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 here. 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 (2017) 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 An Example

This section develops a three-period example that illustrates some of the main ideas. The worker's constant ability is $\theta \sim \mathcal{N}(0, \Sigma)$ and unknown to all. In round $t = 1, 2, 3$, the worker generates output equal to

$$X_t = A_t + \theta + \varepsilon_t, \tag{1}$$

with $\varepsilon_t \sim \mathcal{N}(0, 1)$ independent across rounds. Here, $A_t \geq 0$ is the worker's effort, a hidden action that entails a per-period cost $c(A) = A^2/2$.

In a standard career concerns model, the output X_t is observed by the market, which uses it to update its belief $\mu_t = \mathbf{E}[\theta \mid (X_s)_{s \leq t}]$ and its expectation A_t^* regarding the worker's effort in round t (a deterministic function of time). In the absence of enforceable contracts, the worker is paid upfront for the value of the expected output, $\mu_t + A_t^*$. Because the worker's effort affects the market's expectation about θ, but A_t^* is fixed, the worker chooses effort to maximize

$$\sum_{t=1}^{3} \delta^{t-1} \mathbf{E}\left[\mu_t - c(A_t)\right]. \tag{2}$$

An equilibrium is obtained when $A = A^*$. Plainly, in any equilibrium, $A_3 = 0$, as period 3 is the last round.

We modify this standard baseline as follows. Past output is not observed by the market (a sequence of short-run employers). Rather, in round t, the market observes a signal or *rating* Y_t of past output levels, and nothing else. Hence, $\mu_t = \mathbf{E}[\theta \mid Y_t]$. The rating is chosen by a rater, who has commitment power and observes output. By the revelation principle, we set Y_t equal to μ_t, the worker's *reputation*.

Here, we focus on linear rating schemes. That is, the rater combines past output realizations using some weights that are common knowledge into some figure that is disclosed to the market. Linearity implies that the market belief μ_t is also normally distributed. By the familiar projection theorem for jointly normal variables,

$$\mu_t := \mathbf{E}[\theta \mid Y_t] = \mathbf{E}[\theta] + \frac{\mathbf{Cov}[\theta, Y_t]}{\mathbf{Var}[Y_t]}(Y_t - \mathbf{E}[Y_t]).$$

Hence, the constraint that the random variables Y_t, μ_t be equal reduces to the scalar condition that $\mathbf{Cov}[\theta, Y_t] = \mathbf{Var}[Y_t]$.

This model encompasses the standard baseline, in which the market observes all past output levels. This is because the belief derived from those observations is a summary statistic. Given a sample of t draws from (1), the posterior belief equals

$$\mu_{t+1} = \frac{\frac{1}{\sigma^2} \sum_{s=1}^{t} (X_s - A_s^*)}{\frac{1}{\Sigma} + \frac{t}{\sigma^2}} = \frac{\Sigma}{1 + t\Sigma} \sum_{s=1}^{t} (X_s - A_s^*), \tag{3}$$

(See DeGroot, Ch. 9.5, Thm. 1). Hence, the baseline model is equivalent to a model with a linear rating, in which the round $t+1$ weights assigned to all earlier rounds are given by $\Sigma/(1 + t\Sigma)$. This special case is referred to as *transparency*.

More generally, with three periods only, a linear rating is summarized by a triple,

$u_{1,2}, u_{1,3}, u_{2,3} \in \mathbf{R}^3$, such that, for $t = 2, 3$,

$$Y_t = \sum_{s<t} u_{s,t} \left(X_s - A_s^* \right). \tag{4}$$

Plugging (4) into (2) given (1), the worker's effort is given by the solution to the first-order conditions

$$A_s = \sum_{t=s+1}^{3} \delta^{t-s} u_{s,t}. \tag{5}$$

The goal of the rater is to maximize total expected output $\sum_{t=1}^{3} \mathbf{E}[X_t]$. Given that ability is exogenous, and ability and effort are perfect substitutes, this is equivalent to maximizing total effort over all three periods. Considering (5), it is immediate that

$$\sum_{s=1}^{3} \mathbf{E}[X_t] = \sum_{s=1}^{3} \sum_{t=s+1}^{3} \delta^{t-s} u_{s,t} = \delta^2 u_{1,3} + \delta u_{1,2} + \delta u_{2,3}.$$

Hence, the rater maximizes, for every t,

$$\sum_{s=1}^{t-1} \delta^{-s} u_{s,t},$$

using control variables $(u_{s,t})$, subject to the constraint that $Y_t = \mu_t$. It is readily verified that the solution is given by $u_{1,2}^* = \frac{\Sigma}{1+\Sigma}$, as well as

$$u_{1,3}^* = \frac{\Sigma}{2(1+2\Sigma)} \left(1 - \frac{\sqrt{2}((1-\delta)\Sigma - \delta)}{\sqrt{1+\delta^2+(1-\delta)^2\Sigma}} \right),$$

and

$$u_{2,3}^* = \frac{\Sigma}{2(1+2\Sigma)} \left(1 + \frac{\sqrt{2}((1-\delta)\Sigma + 1)}{\sqrt{1+\delta^2+(1-\delta)^2\Sigma}} \right).$$

These formulas are to be contrasted with those under transparency (the baseline model), in which, as discussed,

$$u_{1,2} = \frac{\Sigma}{1+\Sigma}, \text{ and } u_{1,3} = u_{2,3} = \frac{\Sigma}{1+2\Sigma}.$$

The first conclusion is that transparency is not optimal. In the last round, the market belief is not as precise as the rater's belief.[9] This obfuscation is not obtained by adding noise but by combining past output realizations in a way that prevents the rater's belief from being backed out from the market belief.[10] Moreover, the market belief cannot be deduced from the rater's belief either! Indeed, one can select two output histories that lead to the same rater's belief but distinct ratings.[11] Hence, the rating is not a Markovian function of the rater's belief—however, in the present context, the rater's belief summarizes the payoff-relevant information, provided that the market has access to neither past ratings nor independent information.

The second conclusion is that higher output does not necessarily imply a higher rating. As an example, consider $\delta = 1/2$ and $\Sigma = 4$. Then, $u^*_{2,1} = 4/5$, and

$$u^*_{3,1} = \frac{2}{9}\left(1 - \sqrt{2}\right) < 0, \text{ and } u^*_{3,2} = \frac{2}{9}\left(2\sqrt{2} + 1\right) > 0.$$

Hence, higher effort (and thus higher output) in the first round leads to a lower rating in the third (but a higher in the second) round. This is because of the worker's impatience (relative to the rater). The demotivating impact of this policy on the first-round effort is dampened by its distant prospect. The benefit is reaped in the second round, as effort in that round is disproportionately captured in the final (and next) rating. Such a policy resembles the practice of past-performance benchmarking, whereby good performance is initially rewarded but handicaps future ratings. Benchmarking is usually attributed to ratcheting (higher performance leads to higher expectations), but ratcheting only arises when the rater cannot commit, unlike here. Furthermore, as we will see, benchmarking does not require the worker to be particularly impatient either, when ability changes over time.

The lack of stationarity, unavoidable with a finite horizon, restricts our analysis: the rating policy depends on the number of rounds that have elapsed, while incentives depend on the number of rounds that remain. To uncover the structural properties of the optimal rating, as well as the qualitative features of its time series, we turn our attention to a stationary ability process in the next section.

[9] In contrast to the last round, the second round is not particularly interesting, as the rater only has access to one output observation on which to form the rating—there is no scope for obfuscation.

[10] To keep the example simple, we have not allowed the rater to add extraneous noise to the rating. It is easy to check that doing so is suboptimal.

[11] For instance, select $X_1 = 1, X_2 = 53$ and $X_1' = 10, X_2' = 44$. Then, $Y_3 \neq Y_3'$, but the public belief is the same in both cases.

3 The Model

Time is continuous and the horizon infinite.[12] There is incomplete information. The worker's ability 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

$$\mathrm{d}\theta_t = -\theta_t \,\mathrm{d}t + \gamma \,\mathrm{d}Z_{0,t}, \tag{6}$$

where Z_0 is an independent standard Brownian motion, and $\gamma > 0$.[13,14] Mean-reversion ensures that the variance of θ remains bounded, independent of the market information, thereby accommodating a large class of information structures. The noise Z_0 ensures that incomplete information persists, and thus, the stationary distribution is nontrivial. The specification of the initial variance ensures that the process is stationary. The non-stationary version of the model is analyzed in Section 4.2.

The worker exerts hidden effort. Given effort A_t, cumulative output $X_t \in \mathbf{R}$ solves

$$\mathrm{d}X_t = (A_t + \theta_t)\,\mathrm{d}t + \sigma \,\mathrm{d}Z_{1,t}, \tag{7}$$

where $X_0 = 0$, $\sigma > 0$, and Z_1 is an independent standard Brownian motion. Alongside a sigma-algebra $\mathcal{R}_0$, output X is the only source of information. We refer to the corresponding filtration as $\mathcal{R}$, where $\mathcal{R}_t$ is (the usual augmentation of) $\mathcal{R}_0 \vee \sigma(\{X_s\}_{s \le t})$. This is the rater's information, and it is also available to the worker. Like the rater, he learns about ability by observing output.[15] On path, his belief coincides with the rater's. Let $\mathcal{A}$ denote the set of worker's strategies.[16]

The rater chooses the market's information—formally, an information structure is

[12]Continuous time has two benefits. First, it gives parsimonious formulas. In the Online Appendix, we define the equivalent discrete-time model and solve for the optimal stationary rating, which converges to the continuous-time optimum. The optimal non-stationary rating can also be derived in discrete time, but the resulting formula would in itself nearly exceed the page limit on the article length of this journal. Second, it allows the use of the martingale representation theorem, an essential ingredient of the proof of Theorem 3.5, which motivates the focus on linear ratings.

[13]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.

[14] The unit rate of mean-reversion is a mere normalization, as is its zero mean. Going from a mean-reversion rate of 1 to ρ requires a simple change of variables. The optimal rating remains well-defined and non-degenerate as $\rho \to 0$.

[15]What the worker observes is irrelevant in the following sense. In the equilibria we study, the worker's effort strategy is deterministic even when he observes both outputs and signals. Hence, it remains optimal if he has access to less information. As a result, directly controlling the information available to the worker (an instrument considered by, for instance, Hansen, 2013) is useless.

[16]That is, the set of bounded processes A that are progressively measurable with respect to $\mathcal{R}$.

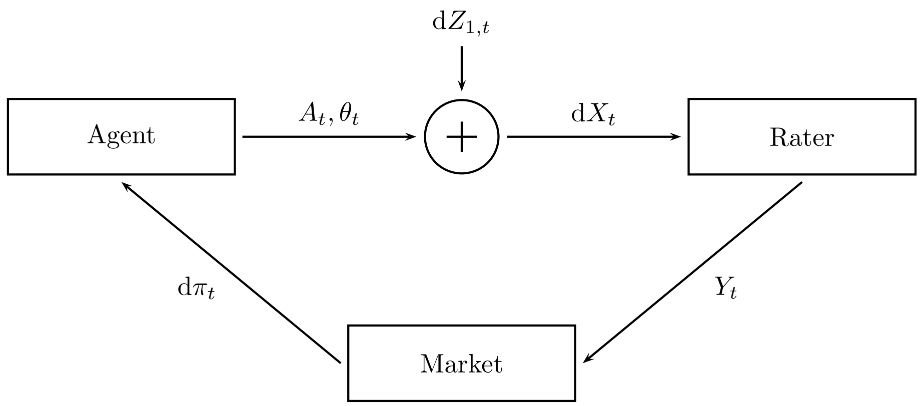

Figure 1: Flow of information and payments between participants.

a collection of sigma-algebras $\mathcal{M}_t \subseteq \mathcal{R}_t$, which need not form a filtration.[17,18] The rater is a designer, not a player. Her objective is irrelevant for now.

Given the cumulative wage π, the market retains[19]

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

whereas the worker receives

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

with $r > 0$ being the common discount rate. The cost $c(\cdot)$ is twice differentiable, with $c'(0) = 0$, $c'' > 0$. Efficiency calls for the (constant) effort A_t that solves $c'(A_t) = 1$. Interactions between agents are summarized in Figure 1.

[17]This is intuitive but heuristic. More formal definitions are provided in Appendix A. Expectations relative to this measure are denoted $\mathbf{E}^*[\cdot]$. This contrasts with the belief P^A of the worker, which captures the actual law of motion. Expectations relative to P^A are denoted $\mathbf{E}[\cdot]$.

[18]We use the star notation when referring to the law on (θ, X) induced by P^{A^*} (see ft. 17). We use the no-star notation when referring to the law on (θ, X) induced by P^A, *i.e.*, the law of motion according to the worker's belief. The market belief is given by the law of (θ, X). Because only the mean ability is payoff-relevant, we call belief the mean ability. The same holds for the worker's belief. We drop the star notation for variance and covariance, for which the distinction is irrelevant.

[19]Formally, π is a continuous, $\mathcal{M}$-adapted process, with $\mathcal{M} := (\mathcal{M}_t)_t$. Note that there is no limited liability, and $\mathrm{d}\pi_t$ can be negative.

Equilibrium has three ingredients. The first is wages: the market is competitive. Given its conjecture A^*, the market pays a flow wage $\mathrm{d}\pi_t$ equal to $\mathbf{E}^*[\mathrm{d}X_t \mid \mathcal{M}_t]$ (possibly negative). Second, the worker chooses A to maximize his payoff. Third, the market has rational expectations: its belief about A is correct.

Definition 3.1 *Fix $\mathcal{M}$. An equilibrium is a profile (A, A^*, π), $A, A^* \in \mathcal{A}$, such that the following hold:*

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

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

2. *(Optimal effort)*

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

3. *(Correct beliefs) It holds that*

$$A^* = A.$$

An equilibrium is called deterministic if A is deterministic, that is, if A is only a function of time.

Because wages depend on the market's perception of ability only, this definition can be simplified. First, the worker's objective can be directly written in terms of this belief. He maximizes his discounted reputation, net of his cost of effort, as formalized below in (8). Second, the market information can be simplified. *A priori*, this information can be very rich—under transparency, for instance, $\mathcal{M}_t$ contains the entire history of output up to time t. However, by the revelation principle, one can work directly with the market belief, fixing a conjectured effort. More precisely, fixing A^*, a sufficient statistic for $\mathcal{M}_t$ is the conditional expectation $\mathbf{E}^*[\theta_t \mid \mathcal{M}_t]$.

Lemma 3.2

(i) Given a wage π that satisfies the zero-profit condition, a deterministic strategy A is optimal if, and only if, it maximizes

$$\mathbf{E} \left[\int_0^\infty e^{-rt} (\mu_t - c(A_t)) \, \mathrm{d}t \right] \tag{8}$$

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

(ii) If (A, A^*, π) is an equilibrium given $\mathcal{M}$, then it is an equilibrium given the information contained in $\mathbf{E}^*[\theta_t \mid \mathcal{M}_t]$.[20]

3.1 Ratings

Given Lemma 3.2(ii), we may restrict attention to information structures generated by some scalar process, or *rating*, Y_t.[21] As mentioned above, we can take this rating to be the market belief, but obviously any non-constant affine transformation of this belief is equivalent to such a rating. In the sequel, uniqueness will be understood in this sense, and we will use whichever equivalent representation is most convenient.

Of particular interest are linear ratings. Formally,

Definition 3.3 *A rating Y is a linear rating if it can be written, for all t,*

$$Y_t = \int_{s \le t} u_{s,t} \, \mathrm{d}X_s,$$

up to some additive constant, for some functions $u_{s,t}$, $s \le t$. It is stationary if for all $s \le t$, $u_{s,t} = u(t-s)$.[22]

That is, a linear rating is one that aggregates observations using some weight function and is stationary if the weight attached to some observation in the current rating is a function of the time that elapsed since that observation occurred, rather than the two calendar times involved. For stationarity to make sense, one needs an appropriate specification of the initial sigma-algebra $\mathcal{R}_0$. We provide details in Appendix A.

In addition, we focus on deterministic equilibria, as defined above. That is, the effort of the worker might depend on calendar time but not on the specific history of outputs or ratings. Owing to the Gaussian structure and the linear payoff structure, such equilibria exist. In addition, with linear ratings, the equilibrium is unique, regardless of whether the rating is stationary.

Proposition 3.4 *Fix a linear rating Y. There exists a unique (deterministic) equilibrium given Y.*[23]

---

[20]Formally, it is an equilibrium given $\{\sigma(\mathbf{E}^*[\theta_t \mid \mathcal{M}_t])\}_t$. Note that, unlike $\mathcal{M}$, $\{\sigma(\mathbf{E}^*[\theta_t \mid \mathcal{M}_t])\}_t$ refers to the market conjecture.

[21]Formally, we can take $\mathcal{M}_t = \sigma(Y_t)$, for some real-valued process Y_t that is progressively measurable with respect to $\mathcal{R}$.

[22]Here and below (Theorem 3.5 and Theorem 3.6), the maps $s \mapsto u_{s,t}$ are integrable and square integrable.

[23]More precisely, here and in what follows, given $\{\sigma(Y_t)\}_t$.

While the focus on deterministic equilibria is standard, it is worth noting that, given that beliefs only enter payoffs via an additive constant, all continuation games are strategically equivalent.[24] Hence, Markov perfection (as defined by Maskin and Tirole, 2001) calls for constant effort (*i.e.*, identical effort after all histories, independent of their length). Accordingly, we start our analysis in Section 4 with stationary ratings, in which case the deterministic equilibrium specifies constant effort.

The definition of a linear rating generalizes that of our example of Section 2. As in the example, the rating is equal to the market belief if, and only if, for all t,

$$\mathbf{Cov}[\theta_t, Y_t] = \mathbf{Var}[Y_t], \tag{9}$$

(and $\mathbf{E}^*[Y_t]$ is normalized to 0). Alternatively, one may verify that a linear rating is equal to the market belief by checking the property that, for all t,

$$\mathbf{E}^*[\theta_t \mid Y_t] = Y_t. \tag{10}$$

Both characterizations are intuitive but require proof. (See Proposition B.1 in Appendix B.) Following information-theoretic terminology, $u(\cdot)$ is the *linear filter* defined by Y. When the filter is a sum of exponentials (*e.g.*, $u(t) = \sum_\ell c_\ell e^{-\delta_\ell t}$), the coefficients (resp., exponents) are the *weights* (resp., *impulse responses*) of the filter.

Aside from tractability, there are practical and conceptual reasons to focus on linear ratings. From a practical perspective, linear ratings encompass a variety of practices. 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 $\mathcal{R}_0$ ensures that this is also well defined

[24] This is not to say that all equilibria involve deterministic effort paths. Such equilibria are not discussed in the career concerns literature. This literature either focuses on the Markov equilibrium (often implicitly, as in Holmström) or considers finite-horizon models, in which case backward induction bites, and there is a unique equilibrium (which is nonstationary Markov). With an infinite horizon, other equilibria certainly exist, here as in Holmström.

for $s \leq 0$.)

Finally, as in our three-period example, transparency can be replicated with a linear rating. Namely, define $\nu_t := \mathbf{E}^*[\theta_t \mid \mathcal{R}_t]$ as the rater's belief. The formula that generalizes (3) to continuous time is given by

$$\nu_t = (\kappa - 1) \int_{s \leq t} e^{-\kappa(t-s)} \left(\mathrm{d}X_s - A_s^* \,\mathrm{d}s\right), \tag{11}$$

where $\kappa := \sqrt{1 + \gamma^2/\sigma^2}$ measures how fast a Bayesian discounts past innovations (output realizations) in her belief, as the state continues evolving. Transparency is a special case of exponential smoothing, with $\delta = \kappa$.

Of course, not all ratings are linear. Linearity rules out coarse ratings, for instance. However, under other rating structures, a deterministic equilibrium (if any) must involve lower effort, as we establish next.

Formally, any deterministic effort level, let alone a constant effort level, can be surpassed by a linear rating.

Theorem 3.5 *If A is a deterministic equilibrium strategy given some rating, then there exists a deterministic equilibrium strategy $\widehat{A}$ given some linear rating, such that, for almost all t, $\widehat{A}_t \geq A_t$.*

This theorem relies on the martingale representation theorem. Its conclusion is very strong, as effort under the linear rating is higher *pointwise*.

There is a second conceptual reason to focus on linear ratings. They are the only ones that deliver a constant quality of information to the market, irrespective of the history. All other schemes have the feature that, in punishing the worker, the rating also punishes (or rewards) the market by changing the quality of the information that it provides. We believe that 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. Retribution should not affect innocent bystanders.

We prove this by demonstrating the equivalence of three notions. Loosely, linear (stationary) ratings are equivalent to Gaussian market beliefs, and Gaussian beliefs are equivalent to deterministic information quality. The statement of the second equivalence is somewhat technical, and we relegate it to Appendix B (see Proposition B.5). The first equivalence plays an important role in our analysis, and so we state it next. Informally, the rating is linear if, and only if, the market belief is normally distributed at all times. That is, a Gaussian belief can be *represented* by a linear rating. Formally, suppose the worker exerts constant effort. A rating is *Gaussian* if, for all $s > 0$, $(Y_t - \mathbf{E}[Y_t], X_t - X_{t-s})$ is jointly stationary and Gaussian, in addition to some regularity properties.[25]

Theorem 3.6 (Representation) *If the worker exerts constant effort and, under such effort, Y is a Gaussian rating, then there exists $u(\cdot)$ (unique up to measure zero sets) such that, up to an additive constant,*

$$Y_t = \int_{s \le t} u(t-s)\,\mathrm{d}X_s. \tag{12}$$

Conversely, given such a function $u(\cdot)$, (12) *uniquely defines a linear rating.*

The decomposition (12) can be interpreted as a regression of Y_t on the continuum of signal increments $\mathrm{d}X_s$, $s \in (-\infty, t]$. It is a new, 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. The proof in Appendix D.2 yields a closed-form for the filter given a Gaussian rating (see also Theorem B.4 in Appendix B.

4 Main Results

The rater has commitment, in the sense that Y is chosen once and for all, and it is common knowledge.[26]

Given the results from the previous section, linear ratings make sense whether the rater strives to maximize output, the quality of the market information, or some function of these two statistics.[27] As the example in Section 2 makes clear, these two objectives are not aligned. Transparency, which can be replicated with the rating that relays the rater's belief to the market, fails to maximize output. They are also not opposed, as providing no information to the market also does not maximize

[25]Formally, for all t, $s > 0$, Y_t is $\mathcal{R}_t$-measurable, and the map $s \mapsto \mathbf{Cov}[Y_t, X_{t-s}]$ is absolutely continuous, with an integrable and square-integrable generalized derivative.

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

[27]Another commonly cited objective of ratings is *stability* over time, as measured by the variance of the belief (Cantor and Mann (2007)). In a Gaussian setting, the stability and quality of market information (as measured by the variance of the type conditional on the belief) are perfect substitutes. This is formally stated in the appendix (see Lemma B.7).

output. A systematic analysis of the trade-off would take us too far afield, but it is straightforward to perform, as we illustrate in Section 5.1.

Here, we focus on output, or equivalent effort maximization, since unlike market information, effort matters for welfare.[28] In this section, we further focus on stationary linear ratings, such that equilibrium effort is constant.

In the case of a stationary linear rating, the formula for equilibrium effort is very simple (see Lemma B.3 in Appendix B). Normalizing the (stationary) variance of the rating to 1, equilibrium effort solves

$$c'(A_t) = \int_0^\infty u(s)e^{-rs}\,\mathbf{Cov}[Y_{t+s},\theta_{t+s}]\,\mathrm{d}s. \tag{13}$$

The right-hand side is the benefit from incremental effort: the rating will increase by a factor $u(s)$ at time $t+s$, and this impacts the market belief and, hence the wage, to the tune of the covariance between the rating and ability. With a stationary rating, this covariance is constant. Given the unit rate of mean-reversion, it equals

$$\mathbf{Cov}[Y_t,\theta_t] = \frac{\gamma^2}{2}\int_0^\infty u(t)e^{-t}\,\mathrm{d}t.$$

Hence, maximizing effort is equivalent to maximizing

$$\left[\int_0^\infty u(t)e^{-rt}\,\mathrm{d}t\right]\left[\int_0^\infty u(t)e^{-t}\,\mathrm{d}t\right], \tag{14}$$

subject to $\mathbf{Var}[Y_t]=1$. Equivalently, normalizing by $\mathbf{Var}[Y_t]$, we maximize

$$\frac{\int_0^\infty u(t)e^{-rt}\,\mathrm{d}t}{\sqrt{\mathbf{Var}[Y]}}\,\mathbf{Corr}[Y,\theta], \tag{15}$$

where $\mathbf{Corr}[Y,\theta]=\mathbf{Cov}[Y,\theta]/\sqrt{\mathbf{Var}[Y]}$ is the correlation between the rating and ability. All formulas derived below are simple applications, and finding the optimal

[28]Under the optimal rating, unlike under transparency as in Holmström, maximizing effort does not always maximize efficiency. When κ is too high, $c'(A)>1$ and equilibrium effort is too high. 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. 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 (adding "white noise," for instance) scales it to any desired lower effort, including the efficient level; see Corollary 4.2. Hence, we identify the range of implementable actions.

rating amounts to finding the function that maximizes this product.

To develop our intuition regarding the optimal rating, let us begin with a simple example: exponential smoothing. Suppose that the rater contemplates selecting a rating from the family

$$Y_t = \int_{s \leq t} e^{-\delta\kappa(t-s)}\, \mathrm{d}X_s, \tag{16}$$

for some $\delta > 0$, which she freely chooses. Using the terminology introduced above, the rater chooses from the family of filters $u(t) = e^{-\delta\kappa t}$. Recall that this family includes transparency (cf. (11)), for $\delta = 1$.

For any other choice, she conceals information from the market. Applying (14), given δ, the worker's effort solves

$$c'(A) = \frac{\frac{1}{r+\delta\kappa}}{\sqrt{\mathbf{Var}[Y]}} \frac{\frac{1}{1+\delta\kappa}}{\sqrt{\mathbf{Var}[Y]}}. \tag{17}$$

Logically, the second term of (17), being the correlation between ability and the rating, is maximized by transparency ($\delta = \kappa$). In that case, it is readily verified that (17) simplifies to

$$c'(A) = \frac{\kappa - 1}{\kappa + r},$$

which characterizes the optimal effort in Holmström.

However, because of the first term of (17), transparency is not optimal (unless $r = 1$, in which case both terms coincide). This first term reflects the worker's time preference. As is the second term, it is single-peaked in δ. Its numerator is decreasing in r, reflecting the fact that a higher rating (as a lower value of δ implies), *if* taken at face value, boosts incentives. However, such a rating also increases the standard deviation of the rating (the denominator). Once the market accounts for such rating inflation, the value of δ that maximizes this term is interior.

The choice of δ that maximizes this first term depends on r: the more patient the worker is, the more a persistent (low δ) rating is effective in conveying incentives (a more patient worker values higher persistence more, and the cost in terms of the increased standard deviation does not depend on his patience). Unless $r = 1$, the optimal choice is not transparency.

Taking derivatives (with respect to δ) in (17) yields as the optimal solution

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

Unsurprisingly, given (17), the rater chooses a rating 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 the worker's preferences, which Bayes' rule ignores. If the worker is patient, it pays to amplify persistence, and δ is low.

To boost effort, the rater would like to calibrate the rating's persistence to the worker's time preference, without necessarily sacrificing correlation. A better correlation calls for transparency and a filter with an exponent κ. Calibrating persistence argues for another exponent.

There are alternatives to exponential smoothing ratings which are equally simple. For instance, moving window ratings are often used. A moving window is given by

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

where $T > 0$ is the window size. But this is a step in the wrong direction: as we prove in appendix, the best system among exponential smoothing ratings yields higher (stationary equilibrium) effort than any moving window rating.[29] In fact, for many parameters, it simultaneously provides higher effort *and* better information (as measured by the variance of the market belief). Smoothing is good both for incentives and information.

Nonetheless, exponential smoothing is not the best rating system. It structure, as summarized by (16), straitjackets the rater. The two competing forces identified in (17) call for a rating with more flexibility.

Theorem 4.1 *The optimal (stationary) rating is unique and given by, for $r \neq \kappa$,*[30,31]

$$u(t) := \frac{1-\sqrt{r}}{\kappa - \sqrt{r}} \sqrt{r} e^{-rt} + e^{-\kappa t}. \tag{18}$$

Hence, the optimal rating is the sum of two exponentials. Note that the impulse response on the second exponential is precisely the one that appears in the rater's own belief ν_t (see (11). The first impulse response reflects the worker's discount rate, r. Not surprisingly, the rater seeks to make the rating more or less persistent according to the worker's time preference.

[29]This is not an artefact of the focus on a single signal, output. In the proof, we allow the rater to condition on an additional source of information.

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

[31]For $r = \kappa$, the optimal rating is proportional to $(1 - (\kappa - \sqrt{\kappa})\, t)\, e^{-\kappa t}$, the rescaled limit of (18). The rating reduces to one exponential term when $r = 1, \kappa^2$. Transparency results for $r = 1$.

Hence, the rating can be written as a sum of two Markov processes, one of which is ν, and the other which we denote I (for "incentive"). That is, for some $\phi \in \mathbf{R}$ (and adjusting for equilibrium effort),

$$Y_t = \phi I_t + (1-\phi)\left(\nu_t + \frac{\kappa+1}{\kappa}A^*\right),$$

where

$$\mathrm{d}I_t = \sqrt{r}\frac{1-\sqrt{r}}{\kappa-\sqrt{r}}\,\mathrm{d}X_t - rI_t\,\mathrm{d}t.$$

This representation has several consequences:

- The optimal rating—viewed as a time series—is not a Markov process (the sum of two Markov processes is not Markov in general, and not here in particular).

- The optimal rating is not a function of the rater's current belief alone (it is not *Markov* in this sense either, *i.e.*, a function of ν only).[32] Put differently, the best way of passing on the worker's performance to future ratings cannot be encapsulated in the impact of this performance on the rater's belief. This might be surprising, given that the rater's belief is the only payoff-relevant variable.

- The optimal rating is a two-state mixture Markov rating—a combination of Markov chains moving at different speeds (Frydman (2005)).

- White noise is harmful: irrelevant noise has no use, as it depresses effort.[33]

These findings echo a large empirical literature documenting various facets or ratings. We have already mentioned the spillover effect in the literature on subjective performance evaluation (see ft. 2). For credit ratings, it is well known that they do not satisfy the Markov property (Altman and Kao (1992), Altman (1998), Nickell, Perraudin, and Varotto (2000), Bangia *et al.* (2002), Lando and Skødeberg (2002), etc.) This has motivated a literature that seeks to produce models of ratings that would replicate the time series. Frydman and Schuermann (2008) find that such a two-state mixture Markov model outperforms the Markov model in explaining credit ratings and also explains economic features of rating data. Other representations

[32]This property is distinct from the first. The rating can be a Markov process without being a function of the belief. The rating can be a function of the belief without being Markov, as in general, if X is Markov, and f is a function defined over X, $f(X)$ is not Markov.

[33]The proof in the appendix allows the rating to be conditioned on white noise and establishes that this information is not used.

are known to capture these data (*e.g.*, as a process with rating momentum; see Stefanescu, Tunaru, and Turnbull (2009)). This is not inconsistent with our result, to the extent that the optimal rating can be described using other variables than I and ν. There is considerable leeway. For instance, the pair (Y, ν) (the rating and the rater's beliefs) leads to a more concrete if less elegant prescription. The rater continues to incorporate some of her private information (via her belief) into the rating. In terms of (Y, ν), Y is a hidden Markov process, with ν being the hidden state. This formulation is used, for instance, by Giampieri, Davis, and Crowder (2005).

From a theoretical viewpoint, what is perhaps most surprising is that two Markov processes suffice to compute the rating. The part of the proof establishing sufficiency, explained in Appendix C, sheds light on this. Viewing the setup as a principal-agent model, 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.

As mentioned, the optimal rating does not use extraneous noise. White noise can only depress effort. Hence, adding the right amount of white noise to a two-state Markov mixture rating achieves whichever desired level of effort. Hence, such rating systems are canonical, independent of the rater's objective in terms of effort.

Corollary 4.2 *Any stationary equilibrium effort level can be achieved by a two-state mixture Markov model plus white noise.*

4.1 Benchmarking

Note that the coefficient on the incentive state of the optimal rating can be negative (see (18)). Hence, better performance might translate into lower future ratings. This occurs whenever $\sqrt{r}$ lies in $[1, \kappa]$. The shape of the optimal rating also depends on the ranking of the impulse responses, r and κ, giving rise to four cases, as illustrated by Figure 2 ($\sqrt{r} \in [0, 1]$, $[1, \sqrt{\kappa}]$, $[\sqrt{\kappa}, \kappa]$, and $[\kappa, \infty)$).

For $(r, \kappa) = (3, 4)$, for instance, the negative weight on $t = 1$ 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 approximately $\tau + .6$). The market accounts for the fact that the rating "understates" performance; the way this is done improves its quality.

To see why better performance can lead to lower ratings, let us focus on $r < \kappa$ and consider selecting a rating from the parameterized family

$$u(t) = de^{-\delta t} + e^{-\kappa t},$$

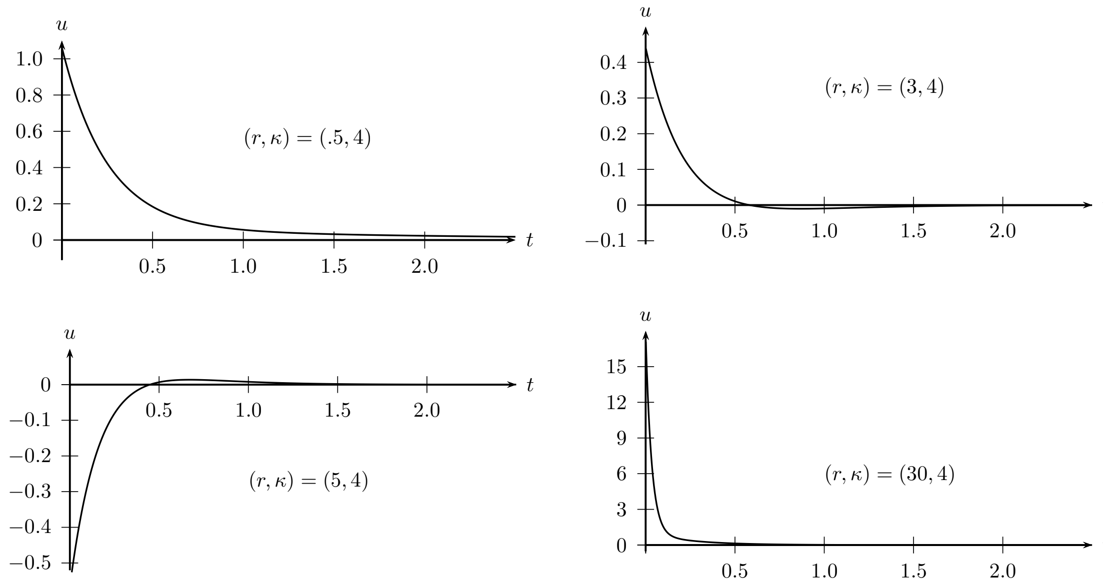

Figure 2: Optimal rating as a function of (r, κ).

where $d \in \mathbf{R}, \delta > 0$ are free. This class is more general than that considered above (*e.g.*, (16)) and embeds the optimal rating, as well as transparency ($d = 0$). Recall from (15) that we can write the marginal cost as a product of two terms, namely:

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

The market is best informed when the rater reveals her own belief. Hence, correlation (Term B) is highest under transparency, setting $d = 0$. Thus, to determine the sign of d, we focus on the first term of (19), Term A. As we have seen in (15), the numerator measures the discounted impact of effort on future ratings, and the denominator corrects for the overall scale of ratings. Roughly, increasing d increases both the impact on future ratings, keeping the scale fixed, and the scale of ratings.

Hence, whether this is desirable depends on which one grows faster. If the worker is very impatient, magnifying the impact via a higher value of d is futile. If mean-reversion is strong, ratings are naturally clustered around the long-run mean, and rating inflation via a high value of d is not a paramount concern. Given that the mean-reversion rate is normalized to 1, the comparison reduces to $r \gtrless 1$.[34]

This benchmarking is a robust phenomenon. Not only does it hold for a broad range of parameters, but it also arises in the extensions we shall consider, when past ratings are observed or additional information is available to the rater or the market. As we saw in Section 2, it also arises with a finite horizon and a fixed ability.

Moreover, it resonates with practice. Murphy (2001) documents the widespread use 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.)

4.2 Non-stationary Setting

Thus far in Section 4, attention has been restricted to stationary ratings. This has required an infinite horizon—indeed, initial conditions that replicated a doubly infinite horizon. To understand the role of stationarity, we consider the polar opposite case here. The horizon is finite, starting at time 0, with no information to begin with, and finishing at time $T < \infty$. The model is otherwise unchanged. Our attention turns to linear (not necessarily stationary) ratings, as introduced in Definition 3.3.

In such a setting, equilibrium effort is not constant. How effort levels at different epochs enter the rater's objective is not innocuous. Maximizing total discounted effort, or maximizing total cumulative effort (given the finite horizon) are obvious candidates. We can solve for both. Here, we focus on the second objective, as the nonstationarity of the optimal rating can then be ascribed to the worker's incentives, rather than the rater's impatience. That is, the rater chooses the rating to maximize

$$\int_0^T A_s \, \mathrm{d}s.$$

In the Online Appendix, Section SA.13, we derive the optimal rating for the case in which the cost is quadratic.

Theorem 4.3 *Suppose $c(A) = A^2$. The optimal rating is unique and given by, for*

[34]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 $r \gtrless 1$.

$r \neq \kappa$ and $0 \leq s \leq t$,

$$u_{s,t} = c_1(t)e^{-r(t-s)} + c_2(t)e^{-\kappa(t-s)} + c_3(t)e^{\kappa(t-s)}, \tag{21}$$

for some functions $c_k(\cdot)$, $k = 1, 2, 3$.

The weights $c_k(\cdot)$ are given in the Online Appendix SA.13. As is clear from (21), the optimal rating is non-stationary. The impulse responses depend on the time elapsed, while the weight on each exponential term depends on calendar time only. Compared to the stationary case, a third exponential branch appears, with exponent κ. However, the rating does not diverge as $t \to \infty$, as the weight $c_3(\cdot)$ vanishes faster than the exponential explodes. Indeed, it is readily verified that:

Corollary 4.4 *As $t \to \infty$, the optimal rating $u_{s,t}$ given by (21), converges to the optimal stationary rating given by (18) (e.g., for all $\tau > 0$, $\lim_{t\to\infty} u_{t-\tau,t} = u(\tau)$).*

Hence, the stationary rating arises as the limit of the optimal rating as the horizon grows large. While calendar time matters, the horizon length T plays no role in the rating.[35] This is because the rating at time t only affects incentives at earlier times, and at such a time $s < t$, different future ratings additively affect incentives. However, the horizon plays a role in determining the optimal effort, as with less time remaining, fewer ratings are affected by current effort. Figure 4.2 illustrates the pattern of effort under the optimal rating for different horizon lengths. As is clear, effort converges pointwise, with a limiting path (as $T \to \infty$) such that effort converges to the stationary level, after an initial phase featuring surprising non-monotonicities.

5 Extensions

Here, we show that the qualitative properties of the results also hold when there are multiple signals available to the rater. In addition, we will allow for situations in which (ii) the market has access to past ratings, (iii) the market also has access to additional signals, and (iv) the worker has multi-dimensional actions.

First, we introduce additional signals. There are good reasons to do so. Workers are evaluated according to a variety of performance measures, both objective and subjective (see Baker, Gibbons, and Murphy (1994)). In the case of a company, in addition to earnings, there is a large variety of indicators of performance (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

[35] That is, T does not appear in (21).

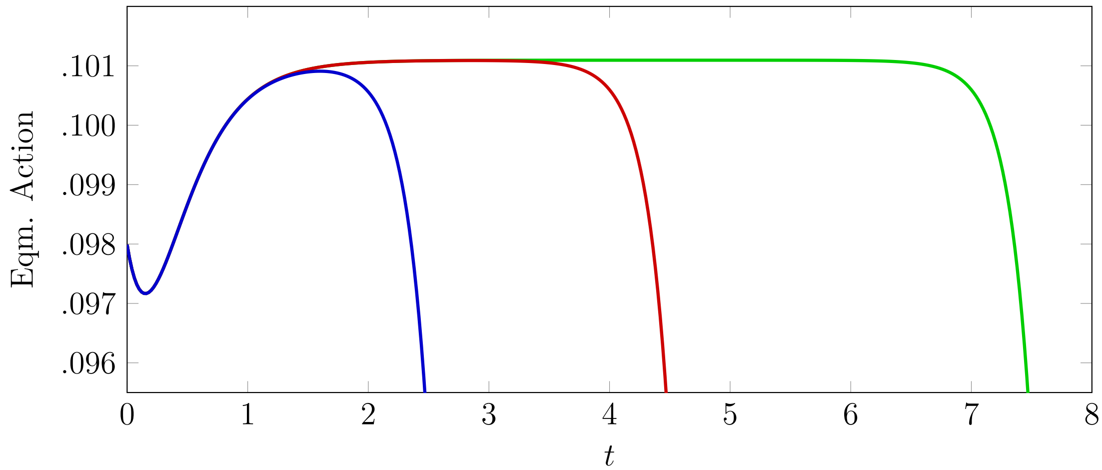

Figure 3: Effort paths for varying horizon lengths ($T = 3, 5, 8$; here, $r = 3, \kappa = \sqrt{2}$).

factors that underlie their ratings (Moody's Investor Services (1991), Moody's Investor Services (1995), Standard & Poor's (1994)).

We model such sources of information as processes $\{S_{k,t}\}$, $k = 2, \ldots, K$ that solve

$$\mathrm{d}S_{k,t} = (\alpha_k A_t + \beta_k \theta_t)\,\mathrm{d}t + \sigma_k\,\mathrm{d}Z_{k,t}, \tag{22}$$

with $S_{k,0} = 0$. Here, $\alpha_k, \beta_k \in \mathbf{R}$, $\sigma_k > 0$, and Z_k is an independent standard Brownian motion. For convenience, we set $S_1 = X$ ($\alpha_1 = \beta_1 = 1$), as output is also a signal. The vector $\mathbf{S} = (S_k)_k$ is observed by the rater and the worker but not by the market. A (stationary) linear rating, then, is a rating that can be written as, for all t,

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

for some maps $\{u_k(\cdot)\}_k$. As with one signal, linear ratings are equivalent to Gaussian ratings, and equilibrium given any linear rating is unique.[36] The coefficient $u_k(s)$ is the weight that the rating at time t assigns to the innovation (the term $\mathrm{d}S_{k,s}$) pertaining to the k-th signal of vintage s.

The next result shows that the optimal rating remains two-state mixture Markov.

[36] Indeed, most results in the appendix are directly established for this more general framework.

Theorem 5.1 *The optimal linear rating is unique and given by*[37]

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

for some constants $\lambda, (d_k^c)_k$.

The constants are given in the Online Appendix, as are all constants appearing in the remainder of this section. Remarkably, the impulse responses of the two exponentials remain the same as with one signal, reflecting transparency (κ) and the worker's discount rate (r). The weights of different signals in the incentive term are ordered according to α_k/β_k. Signals that boost career concerns should see their weight amplified, while those that stifle career concerns should be muted.

Hence, the rating not only merges past and current values of a given signal but also combines the signals of a given epoch, in a fixed proportion. To understand why this second type of useful information is necessary, consider the following extreme example. Suppose, contrary to our maintained convention, that output is only a function of effort. A second signal gives noisy information about ability. That is,

$$\mathrm{d}X_t = A_t\,\mathrm{d}t + \mathrm{d}Z_{1,t}, \text{ and } \mathrm{d}S_{2,t} = \theta_t\,\mathrm{d}t + \mathrm{d}Z_{2,t}.$$

If the rater were to disclose all her information, namely $\{X_s, S_{2,s}\}_{s\leq t}$, then effort $\{A_s\}_{s\leq t}$ would not affect the market belief μ_t, which only depends on $\{S_{2,s}\}_{s\leq t}$. As a result, the worker never puts in any effort. If instead the market only saw the sum

$$\mathrm{d}(X_t + S_{2,t}) = (A_t + \theta_t)\,\mathrm{d}t + \mathrm{d}Z_{1,t} + \mathrm{d}Z_{2,t}, \tag{23}$$

career concerns would arise and effort be positive. This is precisely the (continuous-time) version of Holmström's model. Of course, the optimal rating given by Theorem 5.1 is more sophisticated than the "naive" summation given by (23), but this summation already illustrates why transparency is not optimal.

5.1 Public Ratings

Assuming that past ratings are hidden to the market might be appropriate in some environments in which such feedback can be kept confidential, but it makes little sense if the market is long-lived. Here, we assume that past ratings are observed by

[37]The constant $\kappa := \sqrt{1+\gamma^2\sum_k \beta_k^2/\sigma_k^2}$ generalizes the earlier one and captures the rate of decay of the rater's belief.

the market. As in the baseline model, by the revelation principle, we focus on ratings that are proportional to the market belief. Formally, we call a rating Y *public* when the current rating Y_t includes as much information on θ_t as does the history of current and past ratings $\{Y_s\}_{s\leq t}$. To mark the distinction, we refer to the ratings considered previously in Section 4 as *confidential* (hence, the superscript c, as opposed to p). This constrains the rater, who can no longer treat independently different epochs. The rater is bound by her past ratings. A rating Y is the belief of a public rating if

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

This is a constraint on random variables. As before, it is more useful to express it analytically. Recall that, normalizing the rating such that $\mathbf{E}^*[Y_t] = 0$, requiring that a confidential rating is equal to the belief it induces is equivalent to imposing that:

$$\mathbf{Cov}[Y_t, \theta_t] = \mathbf{Var}[Y_t]. \tag{24}$$

This is not enough to ensure that such a rating remains a belief once past ratings are observed. The problem is that the successive confidential ratings do not exhibit the right autocorrelation. An outside observer who, unbeknown to the rater, would observe the time series of confidential ratings would note that its autocorrelation is "off" and draw inferences from it. Imposing the correct autocorrelation is precisely what is needed. Namely, a rating that satisfies (24) is a public rating, if, and only if, it holds that, for all t and all $\tau \geq 0$,

$$\mathbf{Corr}[Y_t, Y_{t+\tau}] = \mathbf{Corr}[\theta_t, \theta_{t+\tau}] \ (= e^{-\tau}). \tag{25}$$

Equipped with this characterization, we show the following:

Theorem 5.2 *The optimal public rating is unique and given by*

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

for some constants λ, d_k^p.

Hence, publicness distorts the impulse response of the incentive term ($\sqrt{r}$ replacing r). This is easy to understand: the rater must satisfy the constraint (25) on the rating's autocorrelation, which must match the unit mean-reversion rate of ability. This distorts the impulse response of the incentive term away from the rater's favorite value (namely, r), toward 1. The geometric mean of these two values is what turns out to be optimal. This impulse response does not suffice to satisfy the autocorrelation constraint, however, and the weight on this term is also distorted. This distortion is extreme when the rater's information is one-dimensional. If there is only one signal, then $d_k^p = 0$, and transparency obtains. The "continuum" of constraints determines the "one-dimensional continuum" of variables and hence the rating.[38]

As for confidential ratings, adding noise to the public rating defined in Theorem 5.2 allows the rater to precisely achieve any lower effort. (As we show, the proof of Corollary 4.2 can be adapted to public ratings.)

It is straightforward to compute and compare performance (effort) and informativeness (variance of the market belief) in the public and confidential cases. We refer the reader to the working paper for specific comparative statics.[39] In terms of comparisons, the marginal cost of effort satisfy

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

while the variances of the market belief satisfy

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

Hence, effort is lower, but the market is better informed given public ratings, confirming a plausible but not foregone conclusion. This raises a natural question: is requiring ratings to be public equivalent to setting standards of accuracy? To answer this, we plot in Figure 4 the solution (maximum marginal cost of effort) to the two problems—confidential and public ratings—*subject to* an additional constraint, a peg on the variance of the market belief. The solution is a rating similar to the unconstrained one; only the weights on the two exponentials vary.

Fixing precision, there is a maximum effort level that can be induced by the rating. (Curves are truncated at this maximum.) Quality and effort are substitutes over some range: transparency does not maximize effort, whether the rating is public or not. Yet, the effort-maximizing rating does not leave the market in the dark.

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

[38]More precisely, transparency obtains if, and only if, α_k/β_k is independent of k for all k such that $(\alpha_k, \beta_k) \neq (0, 0)$.

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

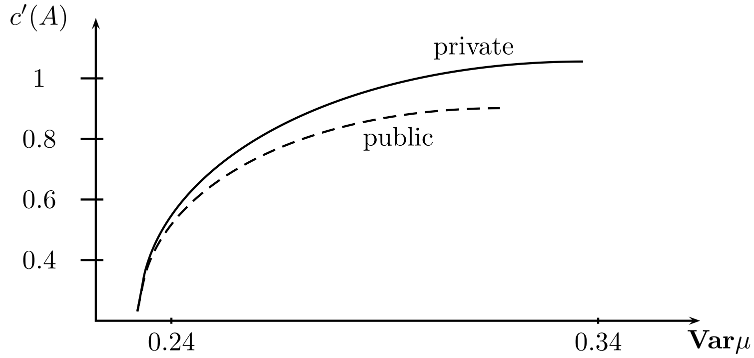

Figure 4: Marginal cost of effort as a function of maximum belief variance, public vs. confidential ratings (here, $(\beta_1, \beta_2, \alpha_1, \alpha_2, \gamma, r, \sigma_1, \sigma_2) = (3, 2, 1/3, 5, 1, 1/5, 1, 2)$).

5.2 Exclusivity

Not all information can be concealed. If the market represents long-run consumers that repeatedly interact with the worker, cumulative output is likely publicly observable. In credit ratings, solicited ratings are based on a combination of information that is widely available to market participants and information that is exclusively accessible to the rater. We refer to this distinction as *exclusive* vs. *non-exclusive* information. The rater 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 rater (we consider the cases of both public and confidential structures, according to whether past information is publicly available).[40] Signals $S_{k,t}$, $k > K_0$, are only observed by the rater and the worker. If $K_0 = 0$, ratings are exclusive, as in Section 4. If $K_0 = K$, we have transparency, and there is nothing for the rater to do. Hence, we focus on $0 < K_0 < K$. As before, we focus on ratings that are proportional to beliefs. We omit the analytic constraints that summarize the restrictions imposed by the rating to be a belief, whether the rating is public or confidential (they are presented in the appendix). Theorem 5.1 also extends to this setup.

Theorem 5.3 *The optimal confidential linear rating is unique*[41] *and given by, for*

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

[41] Under non-exclusivity (in particular, under transparency), the parameters of the publicly available signals can be such that perverse incentives arise and exerting the lowest possible effort

$k \leq K_0$, $u_k = 0$ and $k > K_0$,

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

for some constants $(d_k^c)_k$.[42]

Theorem 5.4 *The optimal non-exclusive linear public rating 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 $k > K_0$,

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

for some constants $d^n, (d_k^e)_k$ and $\delta > 0$.[43]

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 rater's belief. As before, it can be restated as a system in which the rater 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 rater does not need to observe the realized values of the non-exclusive signals to incentivize the worker.[44] However, non-exclusivity affects the quality of

maximizes reputation (see the appendix for sufficient conditions that rule this out). If effort under the rating given here is nil, then effort is nil under any rating.

[42]With non-exclusive ratings, it is a matter of convention whether the belief conveyed by the rating already incorporates the information freely available to the market. In the *interim* case, the belief based solely on the information communicated by the rater must be combined with the non-exclusive signals into a posterior belief (in which case, $u_k = 0$ for $k \leq K_0$). We attempt to preserve as much as possible the analogy with the solution in the exclusive case. This calls for the interim approach for confidential ratings and the alternative, *posterior* approach for public ratings.

[43]The parameters d^n, d_k^e are elementary functions of δ. The parameter δ is the only parameter in the paper that is not solved for in closed-form in the appendix. In fact, we prove that it *cannot* be solved for. It is a root of an irreducible polynomial of degree 6. Using Galois theory, we show that it cannot be expressed in terms of radicals. However, we show that the polynomial admits two positive roots, and we indicate how to select the correct one.

[44]This is not obvious from the statement of Theorem 5.4 because we chose to state the optimal non-exclusive public rating as a posterior belief.

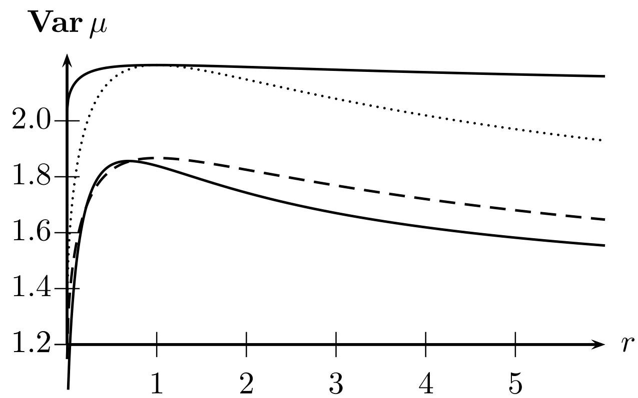

Figure 5: Belief variances (here, $K = 2$ and $(\alpha_k, \beta_k, \sigma_k, \gamma) = (1, 1, 1, 4)$, $k = 1, 2$).

the information available to the market. As an example, consider Figure 5, which 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 about 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 rater cannot conceal. 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 for the case of one signal only (dashed line). For nearly all discount rates, however, the presence of non-exclusive information depresses the rater's willingness to disclose information regarding her unshared signal—free information and the information conveyed by the rating are then strategic substitutes.[45]

5.3 Multiple Actions

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

[45]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. 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.[46] In credit rating, 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).

Our model can accommodate such concerns. We illustrate how in the context of confidential exclusive ratings. We allow for multiple signals, as the main issue is to understand which signals the rating should stress, given the multiplicity of tasks.

Suppose that the worker simultaneously engages in L tasks A_ℓ, with a total cost of effort that is additively separable across these tasks, given by $\sum_{\ell=1}^{L} c(A_\ell)$.[47] For concreteness, assume that $c(A_\ell) = cA_\ell^2$, $c > 0$, although the method applies more generally. Signals are given by, for all $k = 1, \dots, K$,

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

with $\sum_\ell \alpha_{1,\ell} \neq 0$. The rater's goal is to maximize output $X = S_1$, as before. A linear rating is characterized by the filter applied to each of the K signals.

First, we show how our results extend to this environment with the correct change of variables. Define a fictitious model with one-dimensional effort A, cost $c(A) = cA^2$, and signals $\widetilde{S}_k$ such that, for all $k = 1, \dots, K$,

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

where

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

Proposition 5.5 *The linear filter of the optimal rating is the same in both the original model and the fictitious model.*

The optimal filters $\{u_k\}_k$ for the fictitious model are given by Theorem 5.1. Hence, using Proposition 5.5, effort in the original model is given by, for $\ell = 1, \dots, L$,

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

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

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

Second, we illustrate why the optimal rating need not discourage effort in unproductive tasks. To make the point most starkly, let us assume that output is only a function of effort A_1; however, the second signal S_2 reflects both effort A_2 and the worker's ability; namely,

$$\mathrm{d}X_t = A_{1,t}\,\mathrm{d}t + \sigma_1\,\mathrm{d}Z_{1,t}, \text{ and } \mathrm{d}S_{2,t} = (A_{2,t} + \theta_t)\,\mathrm{d}t + \sigma_2\,\mathrm{d}Z_{2,t}.$$

Absent any rating, if either only the first signal or both signals are observed, the unique equilibrium involves $A_1 = 0$. This is because action A_1 does not affect learning about ability. For this effort to be positive, it must matter for his reputation and, thus, for the rating. However, since a higher rating only increases the perceived ability if it increases with high values of the second signal, the rating cannot prevent unproductive effort along the second dimension. Indeed, the optimal filters are

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

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,

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

Unproductive effort is the price to pay for productive effort. The point that this example makes is that, whenever productive effort increases with reputation, other tasks become indirectly useful, to the extent that they invigorate career concerns. If the unproductive task is a better channel to manipulate perceived ability than the productive task, the rating should account for both, without disclosing the breakdown. Allowing workers to play World of Warcraft might encourage them to work harder.[48]

6 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.

[48]See Brown and Thomas (2008) on the benefits of online gaming for organizations.

However, 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 (see also Rodina (2016)). Building on Cisternas (2017b), for instance, one might hope to relax this assumption. Given risk neutrality, effort is a yardstick for efficiency. Allowing for CARA preferences would be useful to evaluate the welfare implications of the ratings informativeness. It should be straightforward to extend our analysis to the case of a multi-dimensional type. This might shed light on the well-documented halo effect in performance evaluation, whereby positive ratings in some dimension impact ratings along other dimensions.

Second, in terms of market structure, we assume a competitive market without commitment and a single worker. 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.

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, E.I. (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.

Bol, J., and D. Smith (2011): "Spillover Effects in Subjective Performance Evaluation: Bias and the Asymmetric Influence of Controllability," *The Accounting Review*, 86, 1213–1230.

Brown, J.S., and D. Thomas (2008): "The Gamer Disposition," *Harvard Business Review*, Feb. 14.

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. (2017a): "Career Concerns and the Nature of Skills," working paper, MIT Sloan.

——(2017b): "Two-Sided Learning and the Ratchet Principle," *The Review of Economic Studies*, forthcoming.

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.

DeGroot, M. (1970): *Optimal Statistical Decisions*, New York, U.S.A.: McGraw-Hill.

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.

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

Ely, J.C. (2017): “Beeps,” *The American Economic Review*, 107, 31–53.

Fong, Y.-F., and J. Li (2017): “Information Revelation in Relational Contracts,” *The Review of Economic Studies*, forthcoming.

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.

Holzbach, R.L. (1978): "Rater bias in performance ratings: Superior, self-, and peer ratings," *Journal of Applied Psychology*, 63, 579–588.

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

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.

Maskin, E., and J. Tirole (2001): "Markov Perfect Equilibrium: I. Observable Actions," *Journal of Economic Theory*, 100, 191–219.

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. (2016): "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.

Rodina, D. (2016): "Information Design and Career Concerns," working paper, Northwestern University.

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 The Model: Complements

This appendix adds some missing formal details. We directly provide these details in the context of the multi-signal version studied in Section 5, which includes the one of Section 3 as a special case.

The formal treatment must address two subtleties.

First, stationarity requires an appropriate specification of the initial sigma-algebra $\mathcal{R}_0$, in a way that "replicates an infinite past." To do so, in the stationary case, we introduce two-sided processes. A *two-sided process* is defined on the entire real line, as opposed to the non-negative real 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 independent standard Brownian motions, and we define two-sided outputs and signals based on these two-sided Brownian motions.

Specifically, for all $t \in \mathbf{R}$, we define

$$\theta_t = e^{-t}\bar{\theta} + \int_0^t e^{-(t-s)}\gamma \, \mathrm{d}Z_{0,s}, \tag{26}$$

where $\bar{\theta} \sim \mathcal{N}(0, \gamma^2/2)$, and Z_0 is a two-sided standard BM. Similarly, let $X = S_1$ and, given the two-sided standard BM Z_k, let S_k be the two-sided process defined by (see ft.17),

$$S_{k,t} = \beta_k \int_0^t \theta_s \, \mathrm{d}s + \sigma_k Z_{k,t}. \tag{27}$$

We let $\{\mathcal{R}_t\}_{t \in \mathbf{R}}$ be the natural augmented filtration generated by $\{S_k\}_k$ over the entire real line, which induces the filtration $\mathcal{R}$ on the non-negative real line. Thus, in the stationary case, $\mathcal{R}_0$ includes non-trivial information about θ, equivalent to what the rater would have learned after observing the output and other signals generated by the worker for an infinite amount of time.

In contrast, in the non-stationary case, θ and S_k are regular (one-sided) processes, defined as (26) and (27) for $t \geq 0$ only, and $\mathcal{R}$ is simply the filtration generated by the output and other signals starting from $t = 0$. In particular, $\mathcal{R}_0$ does not include any information about θ.

In both cases, we denote by P^0 the associated probability measure over the worker's ability, output and signal processes.

The second subtlety is standard in continuous-time principal-agent models. To avoid circularity problems where actions depend on the process that they define, (7) and (22) are to be interpreted in the weak formulation of stochastic differential equations (SDE), where Z is a BM (independent standard Brownian motion) that generally depends on A. Here, signal processes are defined for a reference effort level (say, zero effort). The worker's actions impact the distribution of these signal processes, but does not impact the definition of the processes themselves.

The worker's actions do not define the signal process itself, which is fixed once and for all *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 rater observes.

B Formal Statements: Complements

In this appendix, we provide formal statements for claims that are either missing informally stated in the main body. We divide them by section. However, whenever useful, we state them in the most general version that is used in the paper (*e.g.*, with multiple signals, as in Section 5).

From Section 3, recall that equilibrium behavior is determined by the reputation of the worker at time t, μ_t, which the market computes based on its information $\mathcal{M}_t$. In its most general and abstract form, the rater supplies information to the market, and so defines the *information structure* of the market, which is a collection of sigma-algebras $\{\mathcal{M}_t\}_{t \geq 0}$. Information may be confidential or public as in Section 5.1, and it may be exclusive or non-exclusive as in Section 5.2. The very definition of each of these cases may constrain the information structure of the market. In the baseline case, in which the rater provides confidential and exclusive information, there is no particular constraint on $\{\mathcal{M}_t\}_{t \geq 0}$. When the rater provides public information, the previous information disclosed by the rater is available to the market. Hence, the information structure of the market is constrained to be a filtration: for every $s \leq t$, $\mathcal{M}_s \subseteq \mathcal{M}_t$. Finally, when the rater provides non-exclusive information with public signals $S_1, \ldots, S_{K_0}$, the market information is constrained to include the information conveyed by the public signals: $\sigma(\{S_{k,s} : s \leq t, k \leq K_0\}) \subseteq \mathcal{M}_t$, for all t.

A given information structure defines the market belief in the obvious way. However, as explained in Section 4, given Lemma 3.2, one need not specify the complete information structure–the market belief itself suffices. Indeed, for each information structure corresponding to each of the four cases (confidential/public, exclusive/non-exclusive), we consider the equivalent environment in which the rater provides a single scalar rating which is linear in the market belief. We say that ratings are public when they are equal to the market belief derived from a public information structure, or a non-constant affine transformation thereof. We say that ratings are non-exclusive when they are the market belief derived from a non-exclusive information structure, or a non-constant affine transformation thereof. To clarify when ratings are public and when they are not, we occasionally qualify ratings as confidential (as opposed to public), for added emphasis. Unless mentioned otherwise, it is assumed that the rater provides exclusive information.

In our statements below, we describe the conditions required for a rating to be a market belief with or without public/non-exclusive constraints.

B.1 Statements of Section 3

We start with the informal claim made in (9) and (10) regarding whether a linear rating is equal to the induced market belief.

Proposition B.1 (Confidential Belief) *A stationary linear rating Y is a belief for a confidential information structure if, and only if, for all t,*

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

Equivalently, Y is a belief for a confidential rating if, and only if, for all t,

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

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

In ft.15, we mentioned that the exact information that the worker receives is irrelevant. The next lemma makes this precise.

Lemma B.2 *Fix an arbitrary filtration $\mathcal{H}$. If instead the worker observes information captured by $\mathcal{H}$ (as opposed to $\mathcal{R}$) there continues to exist a unique equilibrium, and the equilibrium effort level does not depend on the chosen filtration $\mathcal{H}$.*

In Section 4.1, the formula for the optimal effort given a linear rating is used in Equation 19. This follows from the next, more general lemma.

Lemma B.3 *Let Y be a linear stationary rating with normalized variance,* $\mathbf{Var}[Y_t] = 1$.[49] *The unique equilibrium effort level A is constant and determined by*

$$c'(A) = \frac{\gamma^2}{2}\left[\sum_{k=1}^{K}\alpha_k \int_0^\infty u_k(t)e^{-rt}\,\mathrm{d}t\right]\left[\sum_{k=1}^{K}\beta_k \int_0^\infty u_k(t)e^{-t}\,\mathrm{d}t\right]. \tag{28}$$

[49]The somewhat unwieldy statement of this constraint in terms of $\{u_k\}_k$ is given in (35) below.

Hence, effort is proportional to the product of two covariances. The first pertains to the worker: 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 normalized variance. Alternatively, without normalization, we may write (28) in a compact way as

$$c'(A^*) = \left[\sum_{k=1}^{K} \alpha_k \int_0^\infty u_k(t) e^{-rt} \, \mathrm{d}t\right] \frac{\mathbf{Cov}[Y_t, \theta_t]}{\mathbf{Var}[Y_t]}. \tag{29}$$

Finally, we state below the general version of the Representation Theorem 3.6 for the multi-signal framework of Section 5.

Theorem B.4 (Representation) *Let Y be a two-sided process such that Y_t is measurable with respect to $\mathcal{R}_t$. Suppose that, when the worker exerts constant effort over time, the following conditions are satisfied:*

1. *for every $s > 0$, $(Y_t - \mathbf{E}[Y_t], \mathbf{S}_t - \mathbf{S}_{t-s})$ is jointly stationary and Gaussian ;*
2. *for all k, $s \mapsto \mathbf{Cov}[Y_t, S_{k,t-s}]$ is absolutely continuous, with integrable and square integrable generalized derivative.*

Then, there exist unique (up to measure zero sets) integrable and square-integrable functions u_k, $k = 1, \ldots, K$, such that, up to an additive constant,

$$Y_t = \sum_{k=1}^{K} \int_{s \le t} u_k(t-s) \, \mathrm{d}S_{k,s}. \tag{30}$$

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

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

This final proposition states that the belief has deterministic (co)variances if, and only if, it is normally distributed. For simplicity, we state the result for the case of output as only signal. The result does not require stationarity, and thus assumes that the initial information $\mathcal{R}_0$ is empty.

Proposition B.5 *Let Y be a (one-sided) rating process progressively measurable w.r.t. $\mathcal{R}$. Suppose that $\mathbf{Cov}\left[Y_t, X_\tau \mid \mathcal{R}_s\right]$ is differentiable in τ, with a derivative that is uniformly Lipschitz continuous in s, and $\mathbf{E}[Y_t^2] < \infty$ and $\mathbf{E}[|Y_t|] < \infty$.*

If the worker exerts a deterministic effort and, under such effort, $(Y_t - \mathbf{E}[Y_t], X_t - X_{t-s})$ is jointly Gaussian in t for $t \geq s$, and both $\mathbf{Cov}\left[Y_t, X_\tau \mid \mathcal{R}_s\right]$ and $\mathbf{Cov}\left[Y_t, \theta_s \mid \mathcal{M}_s\right]$ are deterministic, *for all $t > \tau > s$,*[50] *then there exists $u_{s,t}$, $s \le t$, where $s \mapsto u_{s,t}$ is integrable and square-integrable, such that, for all t,*

$$Y_t = \int_0^t u_{s,t}\,\mathrm{d}X_s.$$

B.2 Statements of Section 4

In the text and ft.29, we claim that exponential smoothing improves on a moving window, even when the rater has access to an extraneous signal and combines it linearly with the information from output. Formally, in the case of exponential smoothing, the rater releases signal

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

with $\delta > 0$, $c \in \mathbf{R}$. With a moving window, the rater releases a signal

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

with $\delta > 0$, $c \in \mathbf{R}$. The *optimal* exponential smoothing (resp., moving window) rating is defined by the choice of (c, δ) (resp., (c, T)) such that (stationary) equilibrium effort is maximized. It holds that:

Proposition B.6 *The optimal confidential exponential smoothing rating yields higher stationary equilibrium effort than any moving window rating.*

As mentioned in ft.27, the precision conveyed by a rating, and its stability over time are perfect substitutes. Formally:

Lemma B.7 *Fix a rating Y. It holds that*

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

B.3 Statements of Section 5

B.3.1 More Signals

Here, we give the explicit functional form of the optimal rating with multiple signals. Define

$$m_\alpha = \sum_{k=1}^{K} \frac{\alpha_k^2}{\sigma_k^2}, \qquad m_{\alpha\beta} = \sum_{k=1}^{K} \frac{\alpha_k \beta_k}{\sigma_k^2}, \qquad m_\beta = \sum_{k=1}^{K} \frac{\beta_k^2}{\sigma_k^2}. \tag{31}$$

We recall that $\kappa := \sqrt{1 + \gamma^2 \sum_k \beta_k^2/\sigma_k^2}$. Also, define

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

[50] That is, non-random functions of (s, t, τ) (resp., (s, t)).

Theorem B.8 *The optimal confidential rating is unique and given by*

$$u_k^c(t) = \frac{\beta_k}{\sigma_k^2}\left(d_k^c \frac{\sqrt{r}}{\lambda} e^{-rt} + e^{-\kappa t}\right), \tag{32}$$

with coefficients

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

Recall that we take ratings as proportional to the market mean belief. For convenience, the formula in (32) implicitly assumes that $\lambda \neq 0$. The proof gives the general formula.

B.3.2 Public Ratings

We start with the counterpart of Proposition B.1 for public beliefs.

Proposition B.9 (Public Belief) *A stationary linear rating Y is a belief for a public information structure if, and only if, it is a belief for a confidential confidential information structure and in addition, for all t and all $\tau \geq 0$,*

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

Equivalently, Y is a belief for a public rating if, and only if, for all t,

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

Here, we give the explicit functional form of the optimal rating with multiple signals.

Theorem B.10 *The optimal public rating 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

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

Recall that λ and d_k^c were introduced and defined in Section B.3.1.

B.3.3 Non-Exclusive Ratings

The following proposition generalizes Propositions B.1 and B.9.

Proposition B.11 (Non-Exclusive Belief) *Let Y be a stationary linear rating. 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,*
$$\mathbf{E}^*[\theta_t \mid \{S_{k,s}\}_{s \leq t, k=1,\ldots,K_0}, \{Y_s\}_{s \leq t}] = Y_t.$$

Equivalently, a rating 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 $\tau \geq 0$,

$$\mathbf{Cov}[S_{k,t}, Y_{t+\tau}] = \mathbf{Cov}[S_{k,t}, \theta_{t+\tau}]. \tag{33}$$

Giving the optimal ratings under non-exclusive ratings requires additional 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 define

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

and

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

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 B.8 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 B.8.

Theorem B.12 *The optimal confidential rating is unique and given by, for $k \leq K_0$, $u_k = 0$, and for signals $k > K_0$,*

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

with coefficients

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

The next theorem gives the optimal rating under public, non-exclusive information.

Theorem B.13 *The optimal non-exclusive public rating 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 of the Online Appendix.

C Overview of the Proofs of Optimal Ratings

This section gives an overview of the proofs that yield the optimal ratings in the four different settings, which correspond to Theorems 4.1 and 5.1 for the baseline case of confidential, exclusive ratings, and to Theorems 5.2, 5.3, and 5.4 for their extensions to the public and/or non-exclusive domains.

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.

The proof of each of the four cases considered (confidential/public and exclusive/non-exclusive) consists of two parts. In the first part, we derive necessary conditions using calculus of variations. These conditions determine a unique candidate for the optimal rating (up to a factor and an additive constant), 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 Theorem (Theorem 3.6, which we prove and establish for the case of multiple signals in Theorem B.4) characterizes all rating in terms of a linear filter $\{u_k\}_k$, which we use as a control variable. Let Y be a rating with normalized variance, $\mathbf{Var}[Y_t] = 1$ (a constraint that can be expressed in terms of $\mathbf{u} := \{u_k\}_k$, see (35) below). The important first step is to determine equilibrium effort, *given* a filter. As Lemma B.3 establishes, the unique equilibrium effort level A is constant and determined by

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

where u_k is defined by Theorem 3.6, given Y. This expresses the equilibrium marginal cost of the worker 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 $\mathbf{u}$:

$$\frac{\gamma^2}{2}\left[\sum_{k=1}^{K}\alpha_k\int_0^\infty u_k(t)e^{-rt}\,\mathrm{d}t\right]\left[\sum_{k=1}^{K}\beta_k\int_0^\infty u_k(t)e^{-t}\,\mathrm{d}t\right]. \tag{34}$$

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 must satisfy. In the simplest case of confidential exclusive ratings, the only constraint is the variance normalization, which is written as follows:

$$\sum_{k=1}^{K} \sigma_k^2 \int_0^\infty u_k(s)^2 \, \mathrm{d}s + \frac{\gamma^2}{2} \int_0^\infty \int_0^\infty U(s)U(t)e^{-|s-t|} \, \mathrm{d}s \, \mathrm{d}t = 1, \tag{35}$$

where $U := \sum_k \beta_k u_k$. The higher dimensionality of the problem is plain in (35). Maximizing (34) subject to (35) 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 (34) and constraint (35) 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.1 of the Online 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 5.2), the objective (34) remains the same, but there are additional constraints on the rating. These capture the restriction that market beliefs are linked to public or non-exclusive ratings. Propositions B.1 and B.9 state these constraints in the exclusive case, and Proposition B.11 does so for the non-exclusive cases. Then, we can apply the Representation Theorem for the case of multiple signals (Theorem B.4) 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.[51] For example, in the public exclusive setting, the constraint (35) is replaced by

$$\sum_{k=1}^{K} \sigma_k^2 \int_0^\infty u_k(t)u_k(t+\tau) \, \mathrm{d}t + \frac{\gamma^2}{2} \int_0^\infty \int_0^\infty U(s)U(t)e^{-|s+\tau-t|} \, \mathrm{d}s \, \mathrm{d}t = 1, \qquad \forall \tau \geq 0. \tag{36}$$

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 just described. For instance, in the

[51]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.

public exclusive setting, we replace (36) by

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

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

where $h(\tau) \coloneqq e^{-r\tau}$. 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)).[52] Instead, we interpret the rater's optimization differently, as a principal-agent model. In this auxiliary model, the agent/worker produces signals and outputs exactly as in the original model and obtains the same payoffs. However, there is no longer a market, nor a rater. Instead, the worker receives transfers from a principal, who observes all outputs and signals, as does the worker. The principal's information at time t is thus $\mathcal{R}_t$, 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.

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 worker chooses an action process A (the agent's strategy) that maximizes, at all t,

$$\mathbf{E}\left[\int_{s\geq t} e^{-r(s-t)}(\mu_s-c(A_s))\,\mathrm{d}s\;\middle|\;\mathcal{R}_t\right].\tag{37}$$

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

[52]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 has a discount rate $\rho \in (0, r)$ and seeks to maximize the *ex ante* payoff

$$\mathbf{E}\left[\int_0^\infty \rho e^{-\rho t}(c'(A_t) + \phi\mu_t(\nu_t - \mu_t))\,\mathrm{d}t\right], \tag{38}$$

where ϕ is some scalar multiplier and $\nu_t = \mathbf{E}[\theta_t \mid \mathcal{R}_t]$, the mean ability of the worker under transparency. The maximization is performed over all strategies A and transfer processes μ such that the action A is incentive compatible, *i.e.*, such that it maximizes (37).

To interpret the principal's objective, it is useful to consider the reward appearing in (38). The term $c'(A_t)$ is the worker's marginal cost, which the rater maximizes in the original model. If the payoff were reduced to this term, the principal might not choose a transfer μ associated with a market belief. However, for the principal-agent and the original model 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, if μ_t is a market belief, then by Proposition B.1, $\mathbf{Cov}[\mu_t, \nu_t] = \mathbf{Var}[\mu_t]$ and $\mathbf{E}[\mu_t] = 0$, which implies $\mathbf{E}\left[\mu_t(\nu_t - \mu_t)\right] = 0$: the second term is equal to zero on average.

If μ is a market belief process associated with a confidential rating, then 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 rater must induce a constant equilibrium effort by the worker. 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 such that (i) effort constant in the limit, (ii) the optimal transfer tends to a market belief, and (iii) the principal's payoff becomes equal to the rater's objective in the original model (the worker's marginal cost). Formally, by sending ρ to 0, the maximum payoff of the principal converges to the conjectured maximum marginal cost from Part I. Because the principal's payoff cannot be lower than the rater'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 transfer 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.

D Proofs

This section includes the proofs for the other two main results (leaving aside the derivation of the optimal rating, sketched in Appendix C). These are Theorem 3.5, establishing that linear ratings cannot be improved upon, as far as deterministic effort is concerned, and Theorem 3.6 (or rather, its generalization to multiple signals, Theorem B.4), showing that Gaussian ratings and linear ratings are equivalent.

All other proofs (including those for the derivation of optimal ratings, sketched in Appendix C) are in the Online Appendix.

D.1 Proof of Theorem 3.5

The proof comprises two steps. The first step lays down first-order conditions on the worker's equilibrium action in terms of a sensitivity parameter associated to the rating. Expressing first-order conditions via a sensitivity parameter (specifically, a Malliavin derivative) is common in the literature on dynamic contracting. Here, however, we must also account for the dependence of output on the worker's ability. The second step uses a martingale representation of the rating and builds on the first-order condition derived in the first step to design a linear rating that induces an effort as least as high as under the original rating.

Step 1. Consider an arbitrary deterministic strategy $A^\dagger$, *i.e.,* a strategy that prescribes an action as a function of time but not as a function of the history of realizations. Let $\mathbf{P}^*$ and $\mathbf{E}^*$ be the probability measure (resp., expectation operator) induced by A^*, and let $\mathbf{P}^\dagger$ and $\mathbf{E}^\dagger$ be the probability measure (resp., expectation operator) induced by $A^\dagger$. Let

$$\begin{aligned}\nu_t^* &= \mathbf{E}^*[\theta_t \mid \mathcal{R}_t],\\ B_t^* &= \frac{1}{\sigma^2}\int_0^t \left(\mathrm{d}X_s - (A_s^* + \nu_s^*)\,\mathrm{d}s\right).\end{aligned}$$

Similarly, let

$$\begin{aligned}\nu_t^\dagger &= \mathbf{E}^\dagger[\theta_t \mid \mathcal{R}_t],\\ B_t^\dagger &= \frac{1}{\sigma^2}\int_0^t \left(\mathrm{d}X_s - (A_s^\dagger + \nu_s^\dagger)\,\mathrm{d}s\right).\end{aligned}$$

By standard arguments, there exists some (deterministic) function $m : \mathbf{R}_+ \to \mathbf{R}$, given by the Riccati equation for linear filters, such that, for all t,

$$\nu_t^* = \int_0^t e^{-\int_s^t (1+m(s'))\,\mathrm{d}s'} m(s)\left(\mathrm{d}X_s - A_s^*\,\mathrm{d}s\right).$$

Alternatively, we may also write ν_t^* as

$$\nu_t^* = \int_0^t e^{-(t-s)} m(s)\,\mathrm{d}B_s^*.$$

An analogous set of equations holds for $\nu^\dagger$, *with the same function* m. Furthermore, by classical results in linear filtering theory, B^* is a standard Brownian motion on $\mathcal{R}$ under $\mathbf{P}^*$, and similarly, $B^\dagger$ is a standard Brownian motion on $\mathcal{R}$ for $\mathbf{P}^\dagger$. Finally, note that

$$X_t = B_t^* + \int_0^t A_s^* \, \mathrm{d}s + \int_0^t \nu_s^* \, \mathrm{d}s = B_t + \int_0^t A_s^* \, \mathrm{d}s + \int_0^t \int_0^i e^{-(i-j)} m(j) \, \mathrm{d}B_j^* \, \mathrm{d}i,$$

and thus B^* and X generate the same information captured by $\mathcal{R}$.

Next, we let

$$V_\infty = \int_0^\infty e^{-rt} \left(\mu_t - c(A_t^*)\right) \mathrm{d}t,$$

and

$$V_t = \mathbf{E}^*[V_\infty \mid \mathcal{R}_t].$$

Observe that V is a $\mathbf{P}^*$-martingale on $\mathcal{R}$, and since X and B^* generate the same information, we can apply the Martingale Representation theorem. We get that there exists a square integrable predictable process ξ on $\mathcal{R}$ such that

$$V_t = V_0 + \int_0^t \sigma e^{-rs} \xi_s \, \mathrm{d}B_s^*, \qquad \text{with } V_0 = -\int_0^\infty e^{-rt} c(A_t^*) \, \mathrm{d}t,$$

where the term σe^{-rs} is a normalization.

We define W_t as the payoff accumulated from time t onwards, assuming the worker follows strategy A^*:

$$W_t = \mathbf{E}^* \left[\int_t^\infty e^{-r(s-t)} \left(\mu_s - c(A_s^*)\right) \mathrm{d}s \,\middle|\, \mathcal{R}_t \right].$$

It is important to note that this continuation payoff does not depend on the actions taken *before* time t.

Suppose A^* is an equilibrium action strategy, thus an optimal plan of action for the worker. At time t, the worker's continuation payoff obtained from A^* is W_t. If instead the worker deviates from the optimal action strategy A^* on the time interval $[t, t+\delta]$, and follows the worker strategy $A^\dagger$ on that interval only, his continuation payoff at time t is

$$W_t' \coloneqq \mathbf{E}^\dagger \left[\int_t^{t+\delta} e^{-r(s-t)} \left(\mu_s - c(A_s^\dagger)\right) \mathrm{d}s \,\middle|\, \mathcal{R}_t \right] + e^{-r\delta} \mathbf{E}^\dagger \left[W_{t+\delta} \mid \mathcal{R}_t \right],$$

where the last term owes to the law of iterated expectations together with the observation that continuation payoffs at a given time are independent of the action path followed before that time.

Note that

$$e^{-r\delta} W_{t+\delta} = W_t + \int_t^{t+\delta} \mathrm{d}\left(e^{-r(s-t)W_s}\right),$$

and, since

$$V_t = \int_0^t e^{-rs} \left(\mu_s - c(A_s^*)\right) \mathrm{d}s + e^{-rt} W_t,$$

we have

$$\mathrm{d}\left(e^{-r(s-t)W_s}\right) = e^{rt}\left(\mathrm{d}V_s - e^{-rs}\mu_s\,\mathrm{d}s + e^{-rs}c(A^*_s)\,\mathrm{d}s\right)$$
$$= e^{-r(s-t)}\left(\sigma\xi_s\,\mathrm{d}B^*_s - \mu_s\,\mathrm{d}s + c(A^*_s)\,\mathrm{d}s\right).$$

Thus,

$$e^{-r\delta}\mathbf{E}^\dagger\left[W_{t+\delta}\mid\mathcal{R}_t\right] = W_t + \mathbf{E}^\dagger\left[\int_t^{t+\delta} e^{-r(s-t)}\left(\sigma\xi_s\,\mathrm{d}B^*_s - \mu_s\,\mathrm{d}s + c(A^*_s)\,\mathrm{d}s\right)\,\middle|\,\mathcal{R}_t\right].$$

We also have the equality

$$\mathrm{d}X_s = \sigma\,\mathrm{d}B^*_s + (A^*_s + \nu^*_s)\,\mathrm{d}s = \sigma\,\mathrm{d}B^\dagger_s + (A^\dagger_s + \nu^\dagger_s)\,\mathrm{d}s,$$

and hence,

$$\sigma\xi_s\,\mathrm{d}B^*_s = \sigma\xi_s\,\mathrm{d}B^\dagger_s + \xi_s(A^\dagger_s - A^*_s)\,\mathrm{d}s + \xi_s(\nu^\dagger_s - \nu^*_s)\,\mathrm{d}s.$$

As $B^\dagger_t$ is a standard Brownian motion on $\mathcal{R}$ under $\mathbf{P}^\dagger$,

$$e^{-r\delta}\mathbf{E}^\dagger\left[W_{t+\delta}\mid\mathcal{R}_t\right]$$
$$= W_t + \mathbf{E}^\dagger\left[\int_t^{t+\delta} e^{-r(s-t)}\left(\xi_s(A^\dagger_s - A^*_s) + \xi_s(\nu^\dagger_s - \nu^*_s) - \mu_s\,\mathrm{d}s + c(A^*_s)\right)\mathrm{d}s\,\middle|\,\mathcal{R}_t\right].$$

Altogether,

$$W'_t = W_t + \mathbf{E}^\dagger\left[\int_t^{t+\delta} e^{-r(s-t)}\left(c(A^*_s) - c(A^\dagger_s) + \xi_s(A^\dagger_s - A^*_s) + \xi_s(\nu^\dagger_s - \nu^*_s)\right)\mathrm{d}s\,\middle|\,\mathcal{R}_t\right].$$

Since A^* is optimal, we must have $W'_t \le W_t$, and thus

$$\mathbf{E}^\dagger\left[\int_t^{t+\delta} e^{-r(s-t)}\left(c(A^*_s) - c(A^\dagger_s) + \xi_s(A^\dagger_s - A^*_s) + \xi_s(\nu^\dagger_s - \nu^*_s)\right)\mathrm{d}s\,\middle|\,\mathcal{R}_t\right] \le 0,$$

for every $A^\dagger$, with an equality if $A^\dagger = A^*$.

We have:

$$\nu^\dagger_s - \nu^*_s = \int_t^s e^{-\int_j^s(1+m(s'))\,\mathrm{d}s'}m(j)\left(A^\dagger_j - A^*_j\right)\mathrm{d}j.$$

Sending δ to 0, it holds that, for almost every t,

$$c(A^*_t) - \xi_t(1+m(t))A^*_t = \sup_a c(a) - \xi_t(1+m(t))a,$$

and hence, $c'(A^*_t) = \xi_t(1+m(t))$, for almost all t. In particular, ξ_t is deterministic, for almost all t.

Step 2. By the Martingale Representation theorem, for every t, there exists a predictable process ϕ_t, adapted to $\mathcal{R}$, such that

$$\mu_t = \int_0^t \phi_{t,s} \, \mathrm{d}B_s^*.$$

Let $\overline{\phi}_{t,s} = \mathbf{E}^*[\phi_{t,s}]$. Consider the confidential linear rating given by

$$Y_t^\star = \int_0^t \overline{\phi}_{t,s} \, \mathrm{d}B_s^*.$$

Let $A^\star$ be a deterministic equilibrium action strategy associated with confidential rating $Y^\star$. Let $\mathbf{P}^\star$ and $\mathbf{E}^\star$ be the induced probability measure and expectation operator, respectively.

The belief process induced by the confidential rating $Y^\star$ is

$$\begin{aligned} \mu_t^\star = \mathbf{E}^\star\left[\theta_t \mid Y_t^\star\right] &= \mathbf{E}^\star\left[\theta_t\right] + \frac{\mathbf{Cov}[\theta_t, Y_t^\star]}{\mathbf{Var}[Y_t^\star]} \left(Y_t^\star - \mathbf{E}^\star\left[Y_t^\star\right]\right) \\ &= K_t + \frac{\mathbf{Cov}[\theta_t, Y_t^\star]}{\mathbf{Var}[Y_t^\star]} Y_t^\star, \end{aligned}$$

for some constant K_t. Note that, as in the main text, we do not specify with respect to which probability law the variance and covariance are to be taken, since, as opposed to expectations, these operators yield the same value no matter the probability law (as long as it is induced by a deterministic action process). Note that K_t does not matter for the worker's incentives. Hence, it is irrelevant in the sequel.

We have

$$\mathbf{Cov}[Y_t^\star, \theta_t] = \mathbf{E}^*[Y_t^\star \theta_t] = \mathbf{E}^*[Y_t^\star \nu_t^*] = \mathbf{E}^*\left[\int_0^t \overline{\phi}_{t,s} \psi_{t,s} \, \mathrm{d}s\right] = \mathbf{E}^*\left[\int_0^t \phi_{t,s} \psi_{t,s} \, \mathrm{d}s\right] = \mathbf{E}^*[\mu_t \nu_t^*].$$

We also have

$$\mathbf{Var}[Y_t^\star] = \mathbf{E}^*[(Y_t^\star)^2] = \int_0^t \overline{\phi}_{t,s}^2 \, \mathrm{d}s = \int_0^t \mathbf{E}^*[\phi_{t,s}]^2 \, \mathrm{d}s \leq \int_0^t \mathbf{E}^*[\phi_{t,s}^2] \, \mathrm{d}s = \mathbf{E}^*[\mu_t^2].$$

As μ is a belief process and A^* is the associated equilibrium action process, $\mu_t = \mathbf{E}^*[\theta_t \mid \mu_t]$, and thus:

$$\mathbf{E}^*[\mu_t(\mu_t - \theta_t)] = \mathbf{E}^*[\mathbf{E}^*[\mu_t(\mu_t - \theta_t) \mid \mu_t]] = \mathbf{E}^*[\mu_t \mathbf{E}[\mu_t - \theta_t \mid \mu_t]] = 0.$$

Hence,

$$\mathbf{E}^*[\mu_t(\mu_t - \nu_t^*)] = \mathbf{E}^*[\mathbf{E}^*[\mu_t(\mu_t - \nu_t^*) \mid \mathcal{R}_t]] = \mathbf{E}^*[\mu_t(\mu_t - \theta_t)] = 0,$$

and $\mathbf{E}^*[\mu_t \nu_t^*] = \mathbf{E}^*[\mu_t^2]$. Combining these observations,

$$\frac{\mathbf{Cov}[Y_t^\star, \theta_t]}{\mathbf{Var}[Y_t^\star]} \geq \frac{\mathbf{E}^*[\mu_t \nu_t^*]}{\mathbf{E}^*[\mu_t^2]} = 1.$$

Let $\xi^\star$ be the analogue of ξ as defined in Step 1 for belief μ, for belief $\mu^\star$.

We immediately get

$$\xi^\star = \frac{\mathbf{Cov}[Y_t^\star, \theta_t]}{\mathbf{Var}[Y_t^\star]}\mathbf{E}^*[\xi_t] \geq \mathbf{E}^*[\xi_t].$$

Recall that, for almost all t, $c'(A_t^*) = \xi_t(1 + m(t))$, and similarly, $c'(A_t^\star) = \xi_t^\star(1 + m(t))$. Thus, $A_t^\star \geq A_t^*$ for almost all t.

D.2 Proof of Theorems 3.6 and B.4

The proof proceeds in two parts. We start with a candidate guess, whose derivation is explained in Section SA.6 of the Online Appendix.[53] 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:

$$C_1^\infty := \frac{\kappa^2 - 1}{2\kappa} \int_0^\infty \bar{f}'(j) e^{-\kappa j} \, \mathrm{d}j, \tag{39}$$

and

$$C_2^\infty := \frac{(\kappa - 1)^2}{2\kappa} \int_0^\infty \bar{f}'(j) e^{-\kappa j} \, \mathrm{d}j. \tag{40}$$

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:

$$u_k(\tau) := C_1^\infty \frac{\beta_k \gamma^2}{\sigma_k^2(\kappa^2 - 1)} e^{\kappa\tau} + C_1^\infty \frac{\beta_k \gamma^2}{\sigma_k^2(\kappa^2 - 1)} e^{-\kappa\tau} - \frac{f_k'(\tau)}{\sigma_k^2} - \frac{\beta_k \gamma^2}{\sigma_k^2 \kappa} \int_0^\tau \sinh(\kappa(\tau - s))\bar{f}'(s) \, \mathrm{d}s, \tag{41}$$

or equivalently,

$$u_k(\tau) := \frac{\beta_k \gamma^2}{\sigma_k^2 \kappa} \left(\frac{\sinh \kappa\tau + \kappa \cosh \kappa\tau}{1 + \kappa} \int_0^\infty e^{-\kappa s} \, \mathrm{d}\bar{f}(s) - \int_0^\tau \sinh \kappa(t - s) \, \mathrm{d}\bar{f}(s) \right) - \frac{f_k'(\tau)}{\sigma_k^2}. \tag{42}$$

Proof of Integrability. We show that every u_k defined by Equation (42) is integrable and square-integrable. To do so, we have to show that

$$(\sinh \kappa t + \kappa \cosh \kappa t) \int_0^\infty e^{-\kappa s} h(s) \, \mathrm{d}s - (1 + \kappa) \int_0^t \sinh \kappa(t - s) h(s) \, \mathrm{d}s \tag{43}$$

[53] 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 $\mathbf{Var}[Y_t \mid \mathcal{R}_{t-s}]$ 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.

is integrable and square-integrable whenever h and h^2 are. We note that (43) 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, (43) is equal to

$$\frac{1}{2}(\kappa+1)\left(e^{\kappa t}\int_t^\infty e^{-\kappa s}h(s)\,\mathrm{d}s + e^{-\kappa t}\int_0^t e^{\kappa s}h(s)\,\mathrm{d}s\right) + \frac{1}{2}(\kappa-1)e^{-\kappa t}\int_0^\infty e^{-\kappa s}h(s)\,\mathrm{d}s. \tag{44}$$

Thus, (43) is non-negative, and showing the integrability of (43) reduces to showing that the integral of (43) converges on $[0, +\infty)$. It is readily verified by differentiation that (43) 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, (43) is integrable. Next, to show that (44) is square-integrable, we show that

$$e^{\kappa t}\int_t^\infty e^{-\kappa s}h(s)\,\mathrm{d}s \tag{45}$$

is square-integrable. As square-integrable functions are closed under additivity, and h is integrable, (45) is the only term of (44) 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,

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

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. Let u_k be defined by (42), or, equivalently, by (41). Let

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

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 $S_{k,t-\tau}$, $\tau \geq 0$, $k = 1, \dots, K$, 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 $\tau \geq 0$ 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 (41),

$$\begin{aligned} g_k'(\tau) = f_k'(\tau) - C_1 \frac{\beta_k\gamma^2}{\kappa^2-1} e^{\kappa\tau} - C_2 \frac{\beta_k\gamma^2}{\kappa^2-1} e^{-\kappa\tau} + \frac{\beta_k\gamma^2}{\kappa} \int_0^\tau \sinh(\kappa(\tau - s)) \bar{f}'(s)\,\mathrm{d}s \\ - \frac{\beta_k\gamma^2}{2} \int_0^\infty U(s) e^{-|\tau-s|}\,\mathrm{d}s. \end{aligned} \tag{46}$$

Further, multiplying (41) 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,

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

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 (39) and (40), 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,

$$\begin{aligned}\int_0^\infty U(s)e^{-|\tau-s|}\,\mathrm{d}s = \frac{2}{\kappa}\int_0^\tau \sinh(\kappa(\tau-j))\bar{f}'(j)\,\mathrm{d}j \\ + C_1^\infty\left[\frac{e^{\kappa\tau}}{\kappa+1}-\frac{e^{\kappa\tau}}{\kappa-1}\right]+C_2^\infty\left[\frac{e^{-\kappa\tau}}{\kappa+1}-\frac{e^{-\kappa\tau}}{\kappa-1}\right].\end{aligned} \tag{47}$$

Plugging the expression of (47) in (46) 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.