



HAL
open science

Motivational Ratings

Johannes Hörner, Nicolas Lambert

► **To cite this version:**

| Johannes Hörner, Nicolas Lambert. Motivational Ratings. 2021. hal-03187510

HAL Id: hal-03187510

<https://hal.science/hal-03187510v1>

Preprint submitted on 1 Apr 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

WORKING PAPERS

N° 1134

August 2020

“Motivational Ratings”

Johannes Hörner and Nicolas S. Lambert

Motivational Ratings*

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

August 26, 2020

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 one of the goals of ratings. For movies or books, it is the main one. For others, however, motivating the rated agent is just

*We would like to thank Sümeýra Akin, Ian Ball, Daria Khromenkova and Konstantin Shamruk 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. Lambert is grateful to Microsoft Research and the Cowles Foundation at Yale University for their hospitality and financial support. Johannes Hörner acknowledges funding from ANR under grant ANR-17-EURE-0010 (Investissements d’Avenir program).

[†]Yale University, 30 Hillhouse Ave., New Haven, CT 06520, USA, and Toulouse School of Economics (CNRS), 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.

as important. Assessments and evaluations concerning teachers, surgeons, managers or even firms are meant to reduce moral hazard.

These two goals are not necessarily aligned. Reputational incentives require the market information to be sensitive to information about the worker’s effort, but career concerns feed off incomplete information: if effort and ability can be perfectly disentangled, career concerns disappear. The purpose of this paper is to examine how the information structure (the “rating”) should be designed to motivate effort. What is the optimal amount of incomplete information?¹ How should different sources of information be combined? At what rate should past observations be discounted?

Our paper addresses these questions within the standard canonical career concerns model à la Holmström—a model without commitment, in which incentives are intrinsically dynamic. The linear-Gaussian structure ensures that, for a broad range of information structures, equilibria in which effort is a deterministic function of time exist. Our attention is restricted to such information structures and equilibria. The goal is to understand how the information structure affects the resulting equilibrium effort. Our first result is a technical stepping stone: career concerns, as measured by effort, are maximized by information structures generated by *linear* ratings. That is, the market is given a one-dimensional rating that is a linear combination of all past performance metrics privately observed by the rater. Linearity refers both to time (aggregation of past and present metrics) and to the kind of information (aggregation of concomitant metrics). This first result already provides two elementary insights: first, adding noise depresses effort and should be avoided;² second, confounding signals of different kinds is always beneficial, as is confounding signals of different epochs.³

The value of confounding derives from a simple trade-off: for the agent to be motivated, it must be that the market wage be sensitive to effort. But the wage is dictated by the market’s belief about ability. Hence, the rating must be informative, and so also reflect the rater’s actual belief about ability. In addition, the optimal rating must depend on both effort and ability in a way that the market cannot invert. In our framework, this is easily seen from our explicit characterization: in the

¹In 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).

²The ineffectiveness of irrelevant conditioning resonates with standard principal-agent theory; see, for instance, Green and Stokey (1983).

³In the literature on subjective performance evaluation, this is known as the spillover effect: the worker’s current-period performance is evaluated in part based on his past performance (*e.g.*, Bol and Smith, (2011)).

stationary case, at any given time, the rating turns out to be a linear combination of the rater’s private belief about the agent’s ability at that time, and another Markovian variable that we call the *incentive state*, which depends on the agent’s preferences.⁴ Specifically, this incentive state depends on the agent’s patience and affects the rate of decay of the rating. If the agent values the far future, this state guarantees that the rating is more persistent than what Bayes’ rule requires; if the agent is nearsighted, this incentive state ensures that the rating swiftly reacts to a positive blip in output, so as to still motivate effort.

Taking a weighted combination of effort and perceived ability (so to speak) does not necessarily mean taking a weighted combination of the signals that the rater has access to. To see this, suppose that the rater has access to two simultaneous signals, both of which are very informative about ability, so that career concerns are weak. Then she might find it best to count the most informative one negatively in the combination. Hence, the resulting rating is a (noisy) convex combination of ability and effort, rather than of the two signals, and puts more weight on effort, relative to ability, than either of these two signals. In this fashion, the market belief becomes more responsive to effort, which boosts career concerns.

Similarly, when the rater decides how to weigh signals of different seniorities in the current rating, she must compare which of these signals is more informative about current ability, and how the current rating matters for effort exerted at the times these signals were obtained. If the ability is very persistent, and the agent very impatient, the older signal is relatively more relevant for ability than for effort than the more recent one (as the current rating only matters for incentives at recent times). Hence, the rater might choose to put negative weight on the older signal, and positive weight on the more recent signal, so as to boost incentives. As a result, the rating system can count past output 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 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)).

What is surprising, perhaps, is that the rating, this linear function of all past private signals of the rater, can be decomposed so simply as the sum of these two Markov state variables—what is known in statistics as a two-state Markov mixture model. (In the nonstationary case, it is a three-state Markov mixture model, with a third state whose impact vanishes over time.) The evolution of the rater’s private belief is given by Bayes rule. The evolution of the incentive state depends on the

⁴We show that this extends to the nonstationary case (with quadratic cost), albeit in a slightly more intricate form involving a third state variable.

worker’s impatience: this state variable discounts past observations at a rate equal to the worker’s discount rate. This is intuitive: when the worker is patient, it is possible to motivate him by the prospect that effort today will affect all future ratings, with weights that are nearly equal. An impatient worker, on the other hand, must be rewarded for effort with a rating that reacts to news promptly, a feature that comes at the expense of informativeness.

Because maximum effort is achieved by such a two-state Markov mixture rating, all constant effort processes that can be induced by some rating can be induced by such rating processes: simply add a noise term to depress effort commensurately. Hence, focusing on such ratings is without loss, independent of the rater’s objective, or her ability to commit. Two-state Markov mixture models are not Markov processes themselves in general. Hence, the time series of optimal ratings fails to satisfy the Markov property.⁵

Obviously, our results rely on the linear-Gaussian framework. However, within this framework, they are robust to the informational environment. In particular, 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 additional nonproprietary information (exclusive vs. nonexclusive ratings). In Sections 5.1 and 5.2, we further discuss how publicness and exclusivity affect the structure of the optimal rating.

As mentioned, our analysis attempts to improve our understanding of rating systems that mitigate moral hazard. Hence, it is not applicable to settings in which the rating’s purpose is to improve the match between consumers and products (movies, books,...). It does not address *rankings* either, in which the goal is to compare alternatives (schools, say), rather than to provide an assessment that has intrinsic meaning. Because we assume a benevolent, committed rater, our analysis is most relevant to professions such as teachers or surgeons, in which ratings are set by a public agency. However, to the extent that we identify all effort levels that can be implemented, and how this can be done, our results are relevant to all environments in which calibrating effort is paramount.

Our analysis builds on the seminal model of Holmström (1999).⁶ A worker exerts

⁵Interestingly, there is significant evidence that, in many cases, observed ratings (based on proprietary rules) cannot be explained by a simple time-homogeneous Markov model. In fact, two-dimensional Markov models have been introduced in the literature precisely to explain observed credit risk dynamics. See, among others, Frydman and Schuermann (2008), or Adomavicius and Tuzhilin (2005). Our model is better suited for understanding doctors’ report cards, or teachers’ evaluations, examples in which the assumption of a benevolent committed designer is more palatable.

⁶Differences between our model and that of Holmström (1999) include the continuous-time setting, mean-reversion in the ability process, and a multidimensional signal structure. See Cisternas

effort that is hidden to the market, which pays him a competitive wage. This wage is equal to the market’s expectation of the worker’s productivity, which combines 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 a rating. 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 maximize the worker’s incentive to exert effort.^{7,8}

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 yields higher output. Furthermore, we establish an equivalence among such ratings, Gaussian ratings, and ratings such that the quality of the market information is deterministic. 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 nonstationary finite-horizon optimal rating (also a mixture Markov model) and show how it converges to the stationary rating.

Simple examples illustrating the main ideas are developed in Section 2. In Sections 3 and 4, we begin with the case of a single signal: output. This allows us to

(2017a) for a specification that is similar to ours in the first two respects.

⁷We shall see that it is easy to depress effort if the maximum effort exceeds the social optimum. As we will see, there is a trade-off between the level of effort and the precision of the market’s information, another legitimate objective for a rating. 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 to 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 of ratings. Credit ratings, for instance, affect the borrowing firm’s default decision (Manso (2013)). Further, ratings may provide information to the worker himself, informing rather than simply motivating his effort choice; see Hansen (2013). Finally, when 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 (2018)).

⁸Throughout, 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.

convey many of the insights. In Section 5, we extend the main results 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.⁹ We show that it can be optimal for the optimal rating system to encourage effort production in dimensions that are unproductive, if this is the only way to also encourage productive effort. The reader interested in the more formal aspects of the analysis is referred to the appendix, which contains not only additional technical ingredients, and some missing 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. When memory is unlimited, similar effects can still be obtained when players’ types change over time, following, for example, a Markov process as in Cole, Dow and English (1995) or Phelan (2006). The added layer of uncertainty increases the difficulty to learn a player’s type. 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

⁹With 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)).

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.

Even closer to our work, Rodina (2017) elegantly shows that transparency maximizes incentives under general conditions and for a large range of technologies, in which the output of the agent need not be a linear Gaussian function of effort and type. Unlike ours, Rodina’s model is based on the one-shot career-concerns model of Dewatripont, Jewitt, and Tirole (1999). Consequently, his analysis cannot 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. Indeed, our analysis relies on the fact that wages only depend on expected output, and so incentives depend on expected ability only, which implies that the trade-off that determines the optimal rating can only arise in a dynamic model.

Finally, in a paper that involves no incomplete information, Fong and Li (2017) show how 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 exhibits those 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. 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 Examples

This section introduces some of the main ideas of this paper in the context of simple examples. A worker supplies output to a market. The market is interpreted as a sequence of short-run employers. The worker lives for T time periods. In what follows, we are mostly interested in the cases $T = 2, 3$.

The worker's per-period output depends on his ability θ . Ability is constant over time, but random and unknown to all, including the worker. It is drawn from a Gaussian distribution with mean 0 and variance Σ , we write $\theta \sim \mathcal{N}(0, \Sigma)$. In round $t = 1, \dots, T$, the worker generates output equal to

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

with $\varepsilon_t \sim \mathcal{N}(0, 1)$, independent of θ and independent across rounds. The unit variance of the noise is for simplicity, and Σ can be interpreted as the uncertainty to noise ratio. Here, $A_t \geq 0$ is the worker's effort, a hidden action that entails a per-period cost $c(A) = A^2/2$. The socially efficient effort level at any round t —that is, the level that maximizes the expected output net of the cost of effort—is $A_t = 1$.

In the standard career concerns model, the market observes output and updates its belief about ability, $\mu_t := \mathbf{E}[\theta \mid (X_s)_{s < t}]$, given its expectation about effort, A_t^* . Attention is restricted to equilibria in which effort (and thus also anticipated effort) is a deterministic function that depends on time, but not on the details of the history. In the absence of enforceable contracts, the worker is paid upfront for the expected value of output, $\mu_t + A_t^*$. From the worker's point of view, A_t^* is fixed, and so he chooses effort to maximize

$$\sum_{t=1}^T \delta^{t-1} \mathbf{E}[\mu_t - c(A_t)], \tag{2}$$

where δ is the worker's discount factor. An equilibrium obtains when $A_t = A_t^*$ for all t . In the last round, effort is nil, as the wage is paid upfront. In earlier rounds, the worker is expected and motivated to manipulate the belief about ability. Given a

sample of t draws from (1), the posterior belief is given by

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

(see, for example, DeGroot (1970), Ch. 9.5, Thm. 1), and $\mu_1 = \mathbf{E}[\theta] = 0$. Outputs from different rounds are equally weighted in this posterior belief. Plugging (3) into (2) given (1) and taking first-order conditions yields as effort¹⁰

$$A_t = \sum_{k=t+1}^T \frac{\delta^{k-t}}{k - 1 + 1/\Sigma}. \quad (4)$$

As is well known, effort increases with uncertainty about ability, Σ . The larger this uncertainty, the bigger the impact of unanticipated effort on the market's posterior belief, and hence, the higher the worker's later wages. Hence, the worker has stronger motivation to work. Depending on Σ and δ , effort might exceed or fall short of first-best. We are primarily interested in the case in which it is inefficiently low (so that the goal of the rating is to increase effort relative to what would occur here). This necessarily obtains when time passes and uncertainty recedes, or in the steady-state that obtains when ability changes over time.

We modify this standard model as follows. The market no longer observes past output. Rather, it observes a *rating* at the beginning of each round. This rating is chosen by a rater. She commits to a rating policy. The rater's goal is to amplify career concerns. That is, she chooses the rating to maximize the sum of efforts, or equivalently the average effort over time. The model is otherwise unchanged—in particular, the wage is equal to expected output.

The rating Y_t is an arbitrary (possibly vector-valued) function of the rater's information. The rater might observe output, or more generally, signals about effort and ability. Let us start with the case it observes two signals S^k , ($k = \text{I, II}$) with

$$S_t^k = \alpha^k A_t + \beta^k \theta + \varepsilon_t^k,$$

where all the noise terms $\varepsilon_t^k \sim \mathcal{N}(0, 1)$ are independent. The coefficients $\alpha^k, \beta^k \in \mathbf{R}$ quantify the relative importance of effort and ability in signal k .

As an example, take $(\alpha^{\text{I}}, \beta^{\text{I}}) = (1, 0)$, and $(\alpha^{\text{II}}, \beta^{\text{II}}) = (0, 1)$, so that signal I only depends on effort, while signal II only depends on ability. If the rater chooses to disclose all its signals, setting $Y_t = (S_s^{\text{I}}, S_s^{\text{II}})_{s < t}$, then only the second component of

¹⁰The details for all computations are in the Online Appendix.

the rating matters for the market belief: $\mu_t = \mathbf{E}[\theta | Y_t] = \mathbf{E}[\theta | (S_s^{\text{II}})_{s < t}]$. Because effort only enters the first component, which is irrelevant for this belief, the worker is no longer motivated to work, and so $A_t = 0$ for all t . *Transparency* fails to achieve efficiency. Yet, by setting $Y_t = S_t^{\text{I}} + S_t^{\text{II}}$, the rater can replicate the baseline outcome in which effort is given by Equation (4), with appropriately redefined noise term.

But the rater can do even better. Suppose it chooses to release a linear combination of S^{I} and S^{II} . How should the weights reflect the characteristics of the two signals? Let $T = 2$, so that we can focus on the rating at the start of round 2,

$$Y_2 = \lambda S_1^{\text{I}} + (1 - \lambda) S_1^{\text{II}}.$$

With our specification,

$$Y_2 = \lambda A_1 + (1 - \lambda)\theta + \lambda \varepsilon^{\text{I}} + (1 - \lambda)\varepsilon^{\text{II}}.$$

Clearly, choosing $\lambda = 0$ or $\lambda = 1$ yields no effort. In the first case, the rating is no longer sensitive to ability, and so the market belief is independent of the rating. In the second case, the rating is no longer sensitive to effort. The optimal rating must reflect the action, for the agent to be motivated to work, and the rating must reflect the ability, for the belief to depend on the rating.

The optimal weight $\lambda \in \mathbf{R}$ is set to maximize A_1 ($A_2 = 0$, as before). This effort solves the first-order condition. Simple algebra yields, for a given λ ,

$$A_1 = \frac{\delta \lambda (1 - \lambda) \Sigma}{(1 - \lambda)^2 \Sigma + \lambda^2 + (1 - \lambda)^2}, \quad (5)$$

so that the optimal weight is given by

$$\lambda = \frac{1}{1 + 1/\sqrt{\Sigma + 1}}. \quad (6)$$

Hence, the weight on the signal about effort increases in the uncertainty regarding ability. Disclosing the sum (or, equivalently, the arithmetic average) of the signals is only optimal when this uncertainty is negligible ($\Sigma \simeq 0$). At the other extreme, as $\Sigma \rightarrow \infty$, $\lambda \rightarrow 1$: in the limit, the rating is only made up of the signal about effort.

The logic is as follows. The larger the weight the rating puts on the signal about

¹¹The resulting effort is given by

$$A_1 = \frac{\delta \Sigma}{2\sqrt{\Sigma + 1}}.$$

effort, the stronger the impact of the worker’s effort on the realized rating. Hence, holding fixed the market’s inference from a given rating, the stronger the worker’s incentive to put in effort, which is desirable. But the market anticipates this: a larger weight on the signal about effort means that the realized rating is less informative about ability. That is, the weight on the signal about effort and uncertainty about ability are substitutes. The more uncertainty, the larger this weight can be without destroying the worker’s incentives—and of course, the larger the equilibrium effort. As the sensitivity of the market’s posterior belief to information about ability grows unbounded, the rating can focus almost exclusively on the signal about effort. The marginal benefit of increasing λ is the increased sensitivity of the rating to effort (at a constant rate); the marginal cost is the decreased sensitivity of the belief to the rating (at a concave rate, given that increasing λ increases the share of the noise in the rating, relative to the ability). Hence, as Σ increases, the optimal weight λ increases.

As a second example, let $\beta^I = \beta^{II} = \alpha^{II} = 1 > \alpha^I$. We may think of the second signal as the output, whereas the first signal puts more weight on ability than effort. Similar computations yield that the optimal weight is equal to $\lambda = \Lambda(\alpha^I)$, with

$$\Lambda(\alpha^I) := \frac{1}{1 - \alpha^I} - \frac{1}{\sqrt{2}} \sqrt{\frac{1 + (\alpha^I)^2}{(1 - \alpha^I)^2} + \Sigma}. \quad (7)$$

Note that

$$\lambda < 0 \quad \Leftrightarrow \quad \Sigma > \frac{1 + \alpha^I}{1 - \alpha^I}.$$

Hence, the weight on one of the signals—here, the first one—can be negative. The logic is the same as before. Indeed, we can rewrite

$$Y_2 = \lambda S_1^I + (1 - \lambda) S_1^{II} = (1 - \lambda(1 - \alpha^I)) A_1 + \theta + \lambda \varepsilon^I + (1 - \lambda) \varepsilon^{II}. \quad (8)$$

Hence, to put arbitrarily more weight on effort than ability requires putting arbitrarily large negative weight on the first signal. With a rating that is a weighted difference between signals, the rater can blunt the impact of ability on the rating, relative to effort, to any degree it desires. And as before, this is desirable when uncertainty increases.

From this example, it may appear that the rater’s ability relies on the availability of multiple contemporaneous signals. What if the rater only observes output? Then it can still improve on the equilibrium outcome that would prevail under transparency—the same as in the baseline model without ratings—by combining observations from

different rounds. For instance, it could release no information at the end of the first round, and release the sum of the first two outputs at the end of the second.

Let $T = 3$, and let us assume here that the rater observes (only) the output, and the rating is *confidential*, so that later employers do not get to see earlier ratings.¹² Hence, the rater must pick both an initial rating Y_2 , released at the start of the second round, and a rating Y_3 at the start of the third. Let us focus on the choice of

$$Y_3 = \lambda X_1 + (1 - \lambda)X_2 = \lambda A_1 + (1 - \lambda)A_2 + \theta + \lambda \varepsilon_1 + (1 - \lambda)\varepsilon_2,$$

a linear combination of the first and second output.¹³ This expression is obviously very similar to (8), and indeed, if the rater wants to maximize overall incentives (the sum of efforts across periods) simple algebra yields that

$$\lambda = \Lambda(\delta),$$

where $\Lambda(\cdot)$ is given by Equation (7). Note that λ can be negative, namely, as above,

$$\lambda < 0 \quad \Leftrightarrow \quad \Sigma > \frac{1 + \delta}{1 - \delta}.$$

The logic is the same as before. When δ is small, effort in the first round is not sensitive to the rating in the last—or at least not as much as the second effort choice—as the worker discounts it more. Hence, viewed from the last round, the first signal is a stronger signal about ability (versus effort), relative to the second signal (in the sense that effort responds to this rating in the second, but hardly in the first round). We are back to the previous example: given one signal that responds to incentives, and another that is mostly about ability, the weight on the latter can become negative as uncertainty increases enough.

We refer to this as *benchmarking*, the practice of counting older performance measures against a worker’s current evaluation. This phenomenon is robust, as we show in Section 4.1. It persists but gets eroded when the ability process is not perfectly persistent. Indeed, if ability evolves over time, the question becomes: how sensitive to the rating are past effort levels (which depends on the worker’s discount rate), versus how informative are past output levels about current ability (which depends on how the change in ability affects the drift in the market’s belief process)? If the worker is relatively more impatient, the older output becomes relatively more informative about ability, and benchmarking arises.

¹²The paper also considers public ratings, which remain observable to the market at all times.

¹³As is readily shown, it is best to set $Y_2 = X_1$, as this is the rater’s only signal at that stage.

For the worker to exert effort, two conditions must be fulfilled. First, the market belief must be responsive to the rating, so that an increase in the rating brings about an increase in the worker’s wage. The more accurate the rating, the more responsive the market belief. Second, given that the worker is impatient, the rating must respond to effort sufficiently fast for him to feel concern. For the worker to be motivated, the rating must react quickly and vigorously, to the tune of the worker’s impatience.

This leads to a trade-off in the design of the rating. Accuracy calls for signals from different rounds to be assigned the same weight in the rating (see Equation (3)), as a way of mitigating the noise clouding each of them. But if the rating assigns the same weight to signals of different maturities, an increase in effort in any given round leads to a small upturn in future ratings only, especially once many rounds have already gone by. This tension can be clearly seen from the formulation with an arbitrary number of rounds T . Namely, assume that the rating in each round is linear in past output, that is, up to an additive constant, Y_t is equal to

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

for some weights u_{st} measuring how output in round s affects the rating in round $t > s$. Then (normalizing without loss the variance of Y_t to 1) the goal of maximizing total effort can be shown to be equivalent to maximizing

$$\sum_{t=1}^T \mathbf{Corr}[\theta, Y_t] \sum_{s=1}^{t-1} \delta^{t-s} u_{s,t}.^{14} \tag{9}$$

The first term is the correlation between rating and ability, and is maximized by transparency. The second captures how impatience affects the allocation of these weights: the more frontloaded they are, the stronger the incentives. This trade-off leads to two competing candidates for the rate at which the rating incorporates signals over time. Indeed, the optimal rating will be seen to be a compromise between them (see Theorem 4.1): it is a *two-state process*—a (non-Markov) process that can be written as the sum of two Markov processes—with summands that decay at the rate dictated by Bayes’ rule and impatience, respectively.

In the context of the last example, this explains why, even when the worker is myopic ($\delta \simeq 0$), so that the last rating matters arbitrarily more for effort in the second round than in the first round, the rating nonetheless assigns nonzero weight to the first observation, as it helps make the rating more informative about ability.

¹⁴This is the analogue of (19) derived in the general model.

As patience increases, the weight λ goes up, and as δ tends to one, the trade-off is resolved, with equal weights assigned to both observations.

Our example has restricted attention to linear ratings—a scalar rating that linearly combines all the signals that the rater has access to. In discrete time, we do not know whether this is without loss of generality. Numerical simulations suggest it might be, but we do not know how to show this. In continuous time, a framework that we adopt for our general analysis, we *prove* that linear ratings are best when it comes to maximizing effort.¹⁵ A second benefit from continuous time is that formulas are considerably more parsimonious. The necessary optimality conditions *within the class of linear ratings* can be solved in discrete time, but the resulting formula does in itself nearly exceed the page limit on the article length of this journal.¹⁶

3 The Model

Time is continuous and the horizon infinite. 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

$$d\theta_t = -\theta_t dt + \gamma dZ_{0,t}, \quad (10)$$

where Z_0 is an independent standard Brownian motion, and $\gamma > 0$.^{17,18} 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 distribution of ability remains nontrivial. The specification of the initial variance ensures that the process is stationary. (A stationarity ability process is used for both the stationary and the nonstationary version of our model studied in Section 4.2.)

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

$$dX_t = (A_t + \theta_t) dt + \sigma dZ_{1,t}, \quad (11)$$

¹⁵Continuous time allows the use of the martingale representation theorem, an essential ingredient of the proof of Theorem 3.5.

¹⁶Although the formula does converge to the continuous-time optimum. See the Mathematica[©] file available from the authors.

¹⁷Throughout, 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.

¹⁸The 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 nondegenerate as $\rho \rightarrow 0$.

where $X_0 = 0$. Here, $\sigma > 0$, and Z_1 is an independent standard Brownian motion. Alongside a sigma-algebra \mathcal{R}_0 , output X is the only source of information. In the nonstationary version of our model, \mathcal{R}_0 is the trivial sigma-algebra and so is an unnecessary addition, but it is useful to have in the stationary version, so that the rater’s information is stationary. We refer to the corresponding filtration as \mathcal{R} , where \mathcal{R}_t is (the usual augmentation of) $\mathcal{R}_0 \vee \sigma(\{X_s\}_{s \leq t})$. This is the rater’s information, and it is also available to the worker. Like the rater, he learns about ability by observing output.¹⁹ On path, his belief coincides with the rater’s. Let \mathcal{A} denote the set of worker’s strategies.²⁰

The rater chooses the market’s information—formally, an information structure is a collection of sigma-algebras $\mathcal{M}_t \subseteq \mathcal{R}_t$, which need not form a filtration.^{21,22} The rater is a designer, a “machine,” not a player. Her objective is irrelevant for now.

Given the cumulative wage π , the market retains²³

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

whereas the worker receives

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

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.

Equilibrium has three ingredients. The first is wages: the market is competitive.

¹⁹What 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.

²⁰That is, the set of bounded processes A that are progressively measurable with respect to \mathcal{R} .

²¹This 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 $\mathbf{E}[\cdot]$.

²²We use the star notation when referring to the law on (θ, X) induced by P^{A*} (see ft. 21). 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.

²³Formally, π is a continuous, \mathcal{M} -adapted process, with $\mathcal{M} := (\mathcal{M}_t)_t$. Note that there is no limited liability, and $d\pi_t$ can be negative.

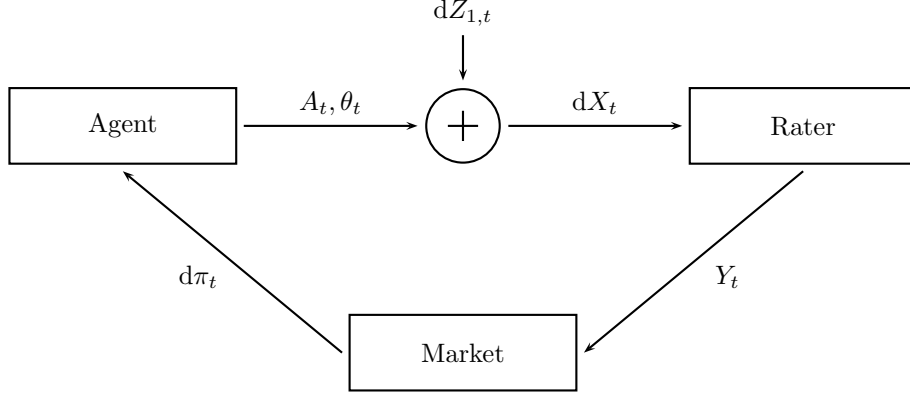


Figure 1: Flow of information and payments between participants.

Given its conjecture A^* , the market pays a flow wage $d\pi_t$ equal to $\mathbf{E}^*[dX_t \mid \mathcal{M}_t]$ (possibly negative). Second, the worker chooses A to maximize his payoff. Third, the market has rational expectations: its conjecture 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 ,

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

2. (Optimal effort)

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

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.

We focus on deterministic equilibria, referred to as equilibria hereafter. The worker's effort can depend on calendar time, but not on the history of outputs or ratings. Here as in other contexts, studying a given effort level is useful for practical applications

and allows for tractability: as we prove, such equilibria admit a simple implementation using linear ratings (plus noise). Implementing a given effort level is a standard focus in contracting papers (see Edmans and Gabaix, 2011; Dittmann and Maug, 2007; Dittman, Maug and Spalt, 2010, which all concentrate exclusively on an effort level, and Grossman and Hart, 1983; Lacker and Weinberg, 1989, which do so predominantly). Deterministic equilibria are also the focus of the career concerns literature.²⁴

Given that beliefs only enter payoffs via an additive constant, all continuation games are strategically equivalent. 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, the analysis in Section 4 starts with stationary ratings, for which the (deterministic) equilibrium involves constant effort. We consider nonstationary equilibria in Section 4.2.

Because wages depend on the market's perception of ability only, Definition 3.1 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 (12). Second, the market information can be streamlined. *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 $\mathbf{E}^*[\theta_t | \mathcal{M}_t]$.

Lemma 3.2

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

$$\mathbf{E} \left[\int_0^\infty e^{-rt} (\mu_t - c(A_t)) dt \right] \quad (12)$$

over \mathcal{A} , where $\mu_t := \mathbf{E}^[\theta_t | \mathcal{M}_t]$ assumes a deterministic conjecture A^* .*

- (ii) *If (A, A^*, π) is an equilibrium given \mathcal{M} , then it is an equilibrium given the information contained in $\mathbf{E}^*[\theta_t | \mathcal{M}_t]$.*²⁵

²⁴This literature typically considers the Markov equilibrium (often implicitly, as in Holmström), which is deterministic, or assumes a finite horizon, in which case the terminal condition yields a unique equilibrium, which is also deterministic. One exception is Cisternas (2017), who studies nondeterministic equilibria when the payoffs of the worker is nonlinear in market beliefs, in which case deterministic equilibria cease to exist.

²⁵Formally, it is an equilibrium given $\{\sigma(\mathbf{E}^*[\theta_t | \mathcal{M}_t])\}_t$. Note that, unlike \mathcal{M} , $\{\sigma(\mathbf{E}^*[\theta_t | \mathcal{M}_t])\}_t$ refers to the market conjecture.

3.1 Ratings

Given Lemma 3.2(ii), we may restrict attention to information structures generated by some scalar process, or *rating*, Y_t .²⁶ As mentioned above, we can take this rating to be the market belief, but obviously any nonconstant 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.

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

$$Y_t = \int_{s \leq t} u_{s,t} dX_s,$$

*up to some additive constant, for some functions $u_{s,t}$, $s \leq t$. It is stationary if for all $s \leq t$, $u_{s,t} = u(t - s)$.*²⁷

That is, a linear rating is one that aggregates observations using some weight function. It 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.

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 ,

$$\mathbf{Cov}[\theta_t, Y_t] = \mathbf{Var}[Y_t], \quad (13)$$

(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 ,

$$\mathbf{E}^*[\theta_t | Y_t] = Y_t. \quad (14)$$

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, that is, $u(t) = \sum_{\ell=1}^L c_\ell e^{-\delta_\ell t}$ for some $L \geq 1$ with $c_\ell, \delta_\ell \in \mathbf{R}$, the coefficients c_ℓ (resp., exponents δ_ℓ) are the *weights* (resp., *impulse responses*) of the filter.

²⁶Formally, 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} .

²⁷Here and below (Theorems 3.5 and 3.6), the maps $s \mapsto u_{s,t}$ are integrable and square integrable.

Of course, not all ratings are linear. Linearity rules out coarse ratings, for instance. However, linear ratings encompass a wide variety of practices. For instance, a *moving window* is a linear rating, with

$$Y_t = \int_{t-T}^t dX_s,$$

for some $T > 0$. (The choice of \mathcal{R}_0 ensures that this is also well defined for $s \leq 0$.) Similarly, *exponential smoothing*, used in some environments, is a linear rating, in which

$$Y_t = \int_{s \leq t} e^{-\delta(t-s)} dX_s,$$

for some choice of $\delta > 0$. A special case is transparency. Namely, define $\nu_t := \mathbf{E}^*[\theta_t | \mathcal{R}_t]$ as the rater's belief. The formula that generalizes (3) to continuous time is

$$\nu_t = (\kappa - 1) \int_{s \leq t} e^{-\kappa(t-s)} (dX_s - A_s^* ds), \quad (15)$$

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. This is exponential smoothing, with $\delta = \kappa$.

Linear ratings enjoy desirable properties. First, from a technical point of view, they ensure uniqueness of the equilibrium, regardless of whether the rating is stationary.

Proposition 3.4 *Fix a linear rating Y . There exists a unique (deterministic) equilibrium given Y .*²⁸

Second, linear ratings are remarkable from the point of view of equilibrium effort and information quality.

In terms of effort, no other rating yields higher effort (whether constant or not). The proof of this result relies on the martingale representation theorem. The result applies to the nonstationary version of the model (i.e., it assumes that the rater has no information at time 0).

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

²⁸More precisely, here and in what follows, given $\{\sigma(Y_t)\}_t$.

Additionally, as explained in the proof of Theorem 3.5, any optimal rating is essentially a linear rating. Adding random noise to the rating has the effect of depressing effort, which allows to flexibly induce effort levels *below* the maximum effort.

In terms of information quality, linear ratings are the only ones that deliver a constant quality of information to the market, irrespective of the history. With any other rating, variations in the rating also punish or reward the market by changing the quality of the information that it provides. We believe that it is a realistic and desirable property for a rating that the history affect the rating level but not the quality of the information: retribution should not affect innocent bystanders. Besides, many market participants have a strong preference for ratings that are not only accurate but also stable. Rating changes have real consequences that are costly to reverse (sending a child to a given school, or picking a surgeon being prime examples).

We relegate a formal statement of this result to the appendix, but conclude this section with a theorem that is instrumental to its proof, and plays a crucial role throughout the analysis. Informally, it establishes an equivalence between the linearity of the rating and the Gaussianity of the market belief. Formally, suppose the worker exerts constant effort. A rating is *Gaussian* if, for all $s > 0$, $(Y_t, X_t - X_{t-s}) - \mathbf{E}[(Y_t, X_t - X_{t-s})]$ is jointly stationary and Gaussian, in addition to some regularity properties.²⁹

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,*

$$Y_t = \int_{s \leq t} u(t-s) dX_s. \quad (16)$$

Conversely, given such a function $u(\cdot)$, (16) uniquely defines a linear rating.

The decomposition (16) can be interpreted as a regression of Y_t on the continuum of signal increments dX_s , $s \in (-\infty, t]$. It is an infinite-dimensional version of the familiar result that a Gaussian variable that depends 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).

²⁹For 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.

4 Main Results

The rater has commitment, in the sense that Y is chosen once and for all, and it is common knowledge. As explained, linear ratings make sense when the rater strives to maximize output or the quality of information.³⁰ As the example in Section 2 makes clear, these two objectives are not aligned. Transparency, which obtains when the rating relays the rater’s belief, fails to maximize output. They are also not opposed, as shunning all information transmission leads to zero effort. 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 equivalently effort maximization, since unlike market information, effort matters for welfare.³¹ 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

$$c'(A_t) = \int_0^\infty u(s)e^{-rs} \mathbf{Cov}[Y_{t+s}, \theta_{t+s}] ds. \quad (17)$$

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

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

³⁰Another 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).

³¹Under the optimal rating, maximizing effort does not always maximize efficiency (unlike in Holmström). When κ is large enough, $c'(A) > 1$ 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), 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, maximizing effort is equivalent to maximizing

$$\left[\int_0^\infty u(t)e^{-rt} dt \right] \left[\int_0^\infty u(t)e^{-t} dt \right], \quad (18)$$

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

$$\frac{\int_0^\infty u(t)e^{-rt} dt}{\sqrt{\mathbf{Var}[Y]}} \mathbf{Corr}[Y, \theta], \quad (19)$$

where $\mathbf{Corr}[Y, \theta] = \mathbf{Cov}[Y, \theta] / \sqrt{\mathbf{Var}[Y]}$ is the correlation between the rating and ability. Finding the optimal rating amounts to maximizing 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

$$Y_t = \int_{s \leq t} e^{-\delta\kappa(t-s)} dX_s, \quad (20)$$

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. (15)), for $\delta = 1$. For any other choice, she conceals information from the market. Applying (18), given δ , the worker's effort solves

$$c'(A) = \frac{1}{r + \delta\kappa} \frac{1}{1 + \delta\kappa} \frac{1}{\sqrt{\mathbf{Var}[Y]}} \frac{1}{\sqrt{\mathbf{Var}[Y]}}. \quad (21)$$

The second term of (21), being the correlation between ability and the rating, is highest under transparency ($\delta = 1$). In that case, it is readily verified that (21) simplifies to

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

which characterizes the optimal effort in Holmström.

However, because of the first term of (21), 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, the more a persistent (low δ) rating delivers incentives (a patient worker values 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 (21) yields as the optimal solution

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

Unsurprisingly, given (21), the rater chooses a rating that is more or less persistent than Bayesian updating according to $r \leq 1$, 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, as mentioned above, moving window systems are often used. But this is a step in the wrong direction: as we prove in appendix, the best rating among exponential smoothing ratings yields higher (stationary) effort than any moving window rating.³² 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.

But exponential smoothing is not the best rating. Its structure straitjackets the rater. The two forces identified in (21) call for a rating with more flexibility.

Theorem 4.1 *The optimal stationary rating is unique and given by, for $r \neq \kappa$,*^{33,34}

$$u(t) := \frac{1 - \sqrt{r}}{\kappa - \sqrt{r}} \sqrt{r} e^{-rt} + e^{-\kappa t}. \quad (22)$$

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

³²This 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.

³³Recall that we take ratings as proportional to the mean market belief. As mentioned, uniqueness is understood as up to such a transformation.

³⁴For $r = \kappa$, the optimal rating is proportional to $(1 - (\kappa - \sqrt{\kappa})t) e^{-\kappa t}$, the rescaled limit of (22). The rating reduces to one exponential term when $r = 1, \kappa^2$. Transparency results for $r = 1$.

own belief ν_t (see (15)). 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.

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

$$dI_t = \sqrt{r} \frac{1 - \sqrt{r}}{\kappa - \sqrt{r}} dX_t - r I_t dt.$$

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).³⁵ 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.³⁶

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. 3). 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

³⁵This 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.

³⁶The proof in the appendix allows the rating to be conditioned on white noise and establishes that this information is not used.

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 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

The coefficient on the incentive state of the optimal rating can be negative (see (22)). Hence, as in the example of Section 2, 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 in 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 $r < \kappa$

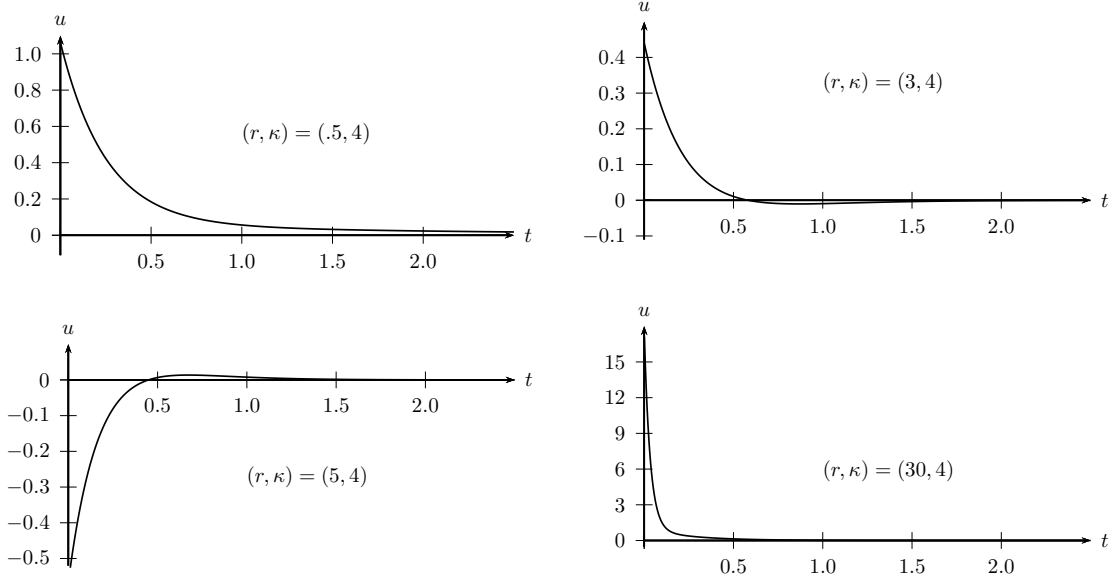


Figure 2: Optimal rating as a function of (r, κ) .

and consider selecting a rating from the parameterized family

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

where $d \in \mathbf{R}$, $\delta > 0$ are free. This class is more general than that considered above (e.g., (20)) and embeds the optimal rating, as well as transparency ($d = 0$). Recall from (19) that we can write the marginal cost as a product of two terms, namely:

$$c'(A) = \frac{\int_0^\infty u(t)e^{-rt} dt}{\sqrt{\mathbf{Var}[Y_t]}} \mathbf{Corr}[Y_t, \theta_t] \propto \underbrace{\frac{\frac{d}{\delta+r} + \frac{1}{\kappa+r}}{\sqrt{\frac{(1+d)^2 + \delta(1+d\kappa/\delta)^2}{1+\delta}}}}_{\text{Term A}} \underbrace{\mathbf{Corr}[Y_t, \theta_t]}_{\text{Term B}}. \quad (23)$$

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 (23), Term A. As we have seen in (19), 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 $r \geq 1$.³⁷

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 Nonstationary 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.³⁸ Here, we focus on the second objective, as in our example of Section 2. Hence, the nonstationarity of the optimal rating can then be ascribed to the worker’s incentives, rather than the rater’s impatience. That is,

³⁷The derivative of Term A evaluated at $d = 0$ is of the same sign as

$$\frac{\kappa + r}{\delta + r} - \frac{\kappa + 1}{\delta + 1}, \tag{24}$$

the sign of which when $\delta < \kappa$ (as when $\delta = r$, its optimal value) is determined by $r \geq 1$.

³⁸Details for the case in which the rater uses some discount factor $R > 0$ are available upon request. The main results generalize: even when $R = r$, the optimal nonstationary rating is a (three-state) mixture Markov model, with effort converging to some long-run effort level that depends on the two discount factors (and coincides with the solution to the stationary model with discount rate $r - R$, which is two-state Markov). Unless κ takes a specific value, transparency does not obtain. Benchmarking results for the same parameters as in Section 4.1, replacing r with $r - R$.

the rater chooses the rating to maximize

$$\int_0^T A_s \, ds.$$

We further assume that the worker has quadratic cost of effort, $c(A) = A^2$. In the Online Appendix, Section [OA.7](#), we derive the optimal rating under generic parameter values.

Theorem 4.3 *The optimal rating is unique and given by, for $r \neq \kappa$ and $0 \leq s \leq t$,*

$$u_{s,t} = c_1(t)e^{-r(t-s)} + c_2(t)e^{-\kappa(t-s)} + c_3(t)e^{\kappa(t-s)}, \quad (25)$$

for some functions $c_k(\cdot)$, $k = 1, 2, 3$ (given in the Online Appendix [OA.7](#)).

Thus, the rating is a three-state mixture Markov model. As is clear from (25), the optimal rating is nonstationary. The impulse responses depend on the time elapsed, while the weight on each exponential term depends on calendar time only. The rating converges as $t \rightarrow \infty$, as the weight $c_3(\cdot)$ vanishes faster than the exponential with impulse response κ explodes. Indeed, it is readily verified that:

Corollary 4.4 *As $t \rightarrow \infty$, the optimal rating $u_{s,t}$ given by (25), converges to the optimal stationary rating given by (22) (e.g., for all $\tau > 0$, $\lim_{t \rightarrow \infty} u_{t-\tau,t} = u(\tau)$).*

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 (that is, T does not appear in (25)). 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 as the marginal cost is linear in effort. However, the horizon affects the optimal effort, as with less time remaining, current effort impact “fewer” ratings. 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 $T \rightarrow \infty$) such that effort converges to the stationary level, after an initial phase featuring surprising nonmonotonicities. Note that the optimal rating and its convergence properties to the optimal stationary rating rely on the assumption that the rater seeks to maximize the average effort of the worker. In general, with other objectives for the rater, and in particular with the rater’s discounting of future effort, the optimal nonstationary rating does not necessarily converge to the optimal stationary rating.

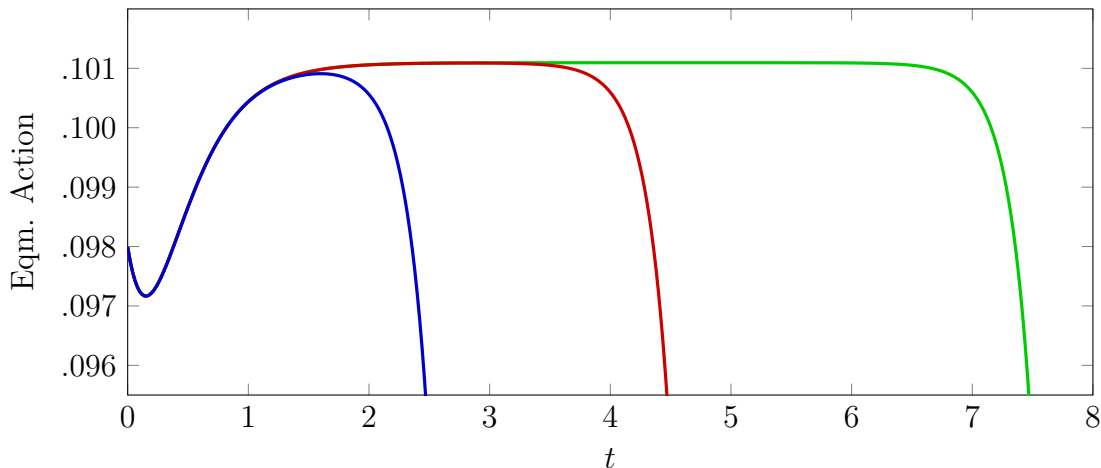


Figure 3: Effort paths for varying horizon lengths ($T = 3, 5, 8$; here, $r = 3, \kappa = \sqrt{2}$).

5 Extensions

Here, we show that our findings extend to the case in which there are multiple signals available to the rater. We also allow for situations in which (ii) the market has access to past ratings, (iii) the market has access to additional signals, and (iv) the worker has multidimensional actions. Throughout, we focus on the stationary case.

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 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 $\{S_{k,t}\}$, $k = 2, \dots, K$ that solve

$$dS_{k,t} = (\alpha_k A_t + \beta_k \theta_t) dt + \sigma_k dZ_{k,t}, \quad (26)$$

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 ,

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

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.³⁹ The coefficient $u_k(s)$ is the weight that the rating at time t assigns to the innovation (the term $dS_{k,s}$) pertaining to the k -th signal of vintage s .

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

Theorem 5.1 *The optimal linear rating is unique and given by⁴⁰*

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

for some constants $\lambda, (d_k^c)_k$.⁴¹

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).

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 see why this second type of obfuscation is useful, consider the following extreme example. Suppose, contrary to our maintained convention, that output depends on effort only. A second signal gives noisy information about ability. That is,

$$dX_t = A_t dt + dZ_{1,t}, \quad \text{and} \quad dS_{2,t} = \theta_t dt + dZ_{2,t}.$$

³⁹Indeed, most results in the appendix are directly established for this more general framework.

⁴⁰The 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.

⁴¹The constant λ is defined by

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

and the constants d_k^c are defined by

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

where

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

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

$$d(X_t + S_{2,t}) = (A_t + \theta_t) dt + dZ_{1,t} + dZ_{2,t}, \quad (27)$$

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 (27), 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 the market. As in the baseline model, and to facilitate the comparison, we focus on ratings that are proportional to the market belief.⁴² In this case, a rating Y is *public* when the current rating Y_t includes as much information on θ_t as does the history of current and past ratings $\{Y_s\}_{s \leq t}$ —which is equivalent to saying that the market observes the history of previous ratings insofar as the information on the current ability is all that matters. 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.

In terms of information precision, the public ratings considered continue to include the extremes, full transparency and zero information. However, “publicness” is a substantial constraint that make public ratings a relatively small subclass of the confidential ratings introduced above. More specifically, a rating Y is the belief of a public rating if

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

⁴²An alternative is to start from any linear rating (not necessarily proportional to market beliefs) and subsequently extract the market belief associated with the observation of past and current ratings. The main difficulty of this approach is the Bayesian updating problem, which to our knowledge is only tractable for more specific forms of rating processes, using a linear filtering methodology. The benefit of the representing the rating directly as a market belief is that no Bayesian updating is necessary, and the conditions that make a rating public can be simply and compactly expressed. However, unlike the baseline case of confidential ratings, it not obvious to us that this restriction to “direct” linear ratings are without loss of generality.

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:

$$\mathbf{Cov}[Y_t, \theta_t] = \mathbf{Var}[Y_t]. \quad (28)$$

This is not enough to ensure that such a rating remains a belief once past ratings are observed. The successive confidential ratings do not exhibit the right autocorrelation. An outside observer who, unbeknown to the rater, would observe the time series of ratings would note that its autocorrelation is “off” and draw inferences from it. Imposing the correct autocorrelation is precisely what is needed. A rating that satisfies (28) is a public rating, if, and only if, it holds that, for all t and all $\tau \geq 0$,

$$\mathbf{Corr}[Y_t, Y_{t+\tau}] = \mathbf{Corr}[\theta_t, \theta_{t+\tau}] (= e^{-\tau}). \quad (29)$$

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 .⁴³

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 (29) 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.⁴⁴

⁴³The constants d_k^p are defined as

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

and the constants κ , λ , and d_k^c are defined as in Theorem 5.1.

⁴⁴More precisely, transparency obtains if, and only if, α_k/β_k is independent of k for all k such that $(\alpha_k, \beta_k) \neq (0, 0)$.

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 in the Online Appendix, 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.⁴⁵ 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. This raises a natural question: is requiring ratings to be public equivalent to setting standards of accuracy? In Figure 4, we plot 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, the rating can induce a maximum effort level (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 the figure makes clear, effort is lower in the public case for any given level of variance. A confidential rating both incentivizes more effort and provides better information than a public one.

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. *nonexclusive*

⁴⁵Which is not to say that these comparative statics are foregone conclusions. For instance, *adding* a signal can lead to a less-informed market.

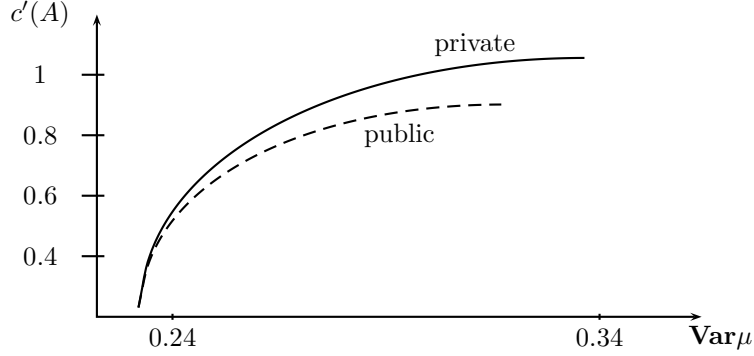


Figure 4: Marginal cost of effort as a function of maximum belief variance, public vs. confidential ratings (here, $(\beta_1, \beta_2, \alpha_1, \alpha_2, \gamma, r, \sigma_1, \sigma_2) = (3, 2, 1/3, 5, 1, 1/5, 1, 2)$).

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 $\{S_{k,s}\}_{s \leq t, k=1, \dots, K_0}$ 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).⁴⁶ 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 nonexclusive linear confidential rating is unique⁴⁷ and given by, for $k \leq K_0$, $u_k = 0$ and $k > K_0$,*

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

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

⁴⁷Under nonexclusivity (in particular, under transparency), the parameters of the publicly available signals can be such that perverse incentives arise and exerting the lowest possible effort 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.

for some constants $(d_k^c)_k$.^{48, 49}

Theorem 5.4 *The optimal nonexclusive linear public rating is unique and given by, for signals $k \leq K_0$,*

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

and for signals $k > K_0$,

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

for some constants d^m , $(d_k^e)_k$ and $\delta > 0$.⁵⁰

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 nonexclusive signals to incentivize the worker.⁵¹ However, nonexclusivity affects the quality of 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 nonexclusive than when it is not (the higher solid vs. the dotted line). However, this is only the case

⁴⁸The constants d_k^c are defined as in Theorem 5.1 with the exception that the value of Δ must be redefined, see Appendix B for details.

⁴⁹With nonexclusive 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 nonexclusive 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.

⁵⁰The parameters d^m , 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. See Appendix B and Section OA.10 of the Online Appendix for details.

⁵¹This is not obvious from the statement of Theorem 5.4 because we chose to state the optimal nonexclusive public rating as a posterior belief.

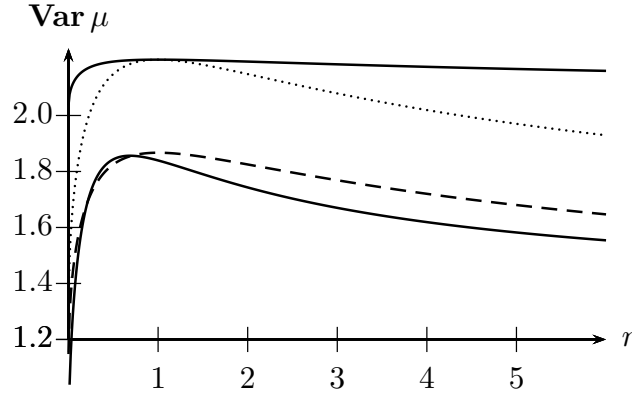


Figure 5: Belief variances (here, $K = 2$ and $(\alpha_k, \beta_k, \sigma_k, \gamma) = (1, 1, 1, 4)$, $k = 1, 2$).

because the market can rely on the nonexclusive signal (the output) in addition to the rating. If a market participant were to rely on the rating alone to draw inferences about ability (lower solid line), he would be worse off under nonexclusivity. This does not imply that the information conveyed by the rating is degraded because of another signal that the rater cannot conceal. As is clear from Figure 5, variance could be even lower if the nonexclusive signal did not exist and we were considering the confidential rating for the case of one signal only (dashed line). For nearly all discount rates, the presence of nonexclusive 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.⁵²

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.

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.⁵³ In credit rating,

⁵²This 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.

⁵³See Porter (2015) for a variety of other examples.

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 $\sum_{\ell=1}^L c(A_\ell)$.⁵⁴ For concreteness, assume that $c(A_\ell) = cA_\ell^2$, $c > 0$, although the method applies more generally. Signals are given by, for all $k = 1, \dots, K$,

$$dS_{k,t} = (\sum_{\ell} \alpha_{k,\ell} A_{\ell,t} + \beta_k \theta_t) dt + \sigma_k dZ_{k,t},$$

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 \tilde{S}_k such that, for all $k = 1, \dots, K$,

$$d\tilde{S}_{k,t} = (\alpha_k A_t + \beta_k \theta_t) dt + \sigma_k dZ_{k,t},$$

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 $\ell = 1, \dots, L$,

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

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,

$$dX_t = A_{1,t} dt + \sigma_1 dZ_{1,t}, \text{ and } dS_{2,t} = (A_{2,t} + \theta_t) dt + \sigma_2 dZ_{2,t}.$$

⁵⁴For a discussion of the restriction implied by separability, see Holmstrom and Milgrom (1991).

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.⁵⁵

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.

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 (2017)). 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

⁵⁵See Brown and Thomas (2008) on the benefits of online gaming for organizations.

to extend our analysis to the case of a multidimensional 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 (2018): “Recommender Systems as Mechanisms for Social Learning,” *Quarterly Journal of Economics*, 133, 871–925.
- 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,” *American Economic Journal: Microeconomics*, forthcoming.
- (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.
- Cole, H.L., J. Dow and W.B. English (1995): “Default, Settlement, and Signalling: Lending Resumption in a Reputational Model of Sovereign Debt,” *International Economic Review*, 36, 365–385.
- 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*, 84, 277–299.
- 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. Holmström (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.
- Holmström, 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.

- Jeulin, T., and M. Yor (1979): “Inégalités de Hardy, semimartingales, et faux-amis,” in *Séminaire de Probabilités XIII 1977/78*, Lecture Notes in Mathematics, ed. by C. Dellacherie, P.-A. Meyer, and M. Weil. Heidelberg: Springer, 332–359.
- 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.
- Phelan, C. (2006): “Public Trust and Government Betrayal,” *Journal of Economic Theory*, 130, 27–43.

- 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 (2017): “Optimal Dynamic Information Provision,” *Games and Economic Behavior*, 104, 329–349.
- Rodina, D. (2017): “Information Design and Career Concerns,” working paper, University of Bonn.
- Sabbagh, L.D. (1969): “Variational Problems with Lags,” *Journal of Optimization Theory and Applications*, 3, 34–51.
- Sannikov, Y. (2008): “A Continuous-Time Version of the Principal-Agent Problem,” *The Review of Economic Studies*, 75, 957–984.
- 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 multisignal version studied in Section 5, which includes the single-signal version 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 nonnegative 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

$$\theta_t = e^{-t\bar{\theta}} + \int_0^t e^{-(t-s)\gamma} dZ_{0,s}, \quad (30)$$

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.21),

$$S_{k,t} = \beta_k \int_0^t \theta_s ds + \sigma_k Z_{k,t}. \quad (31)$$

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 nonnegative real line. Thus, in the stationary case, \mathcal{R}_0 includes nontrivial 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 nonstationary case, θ and S_k are regular (one-sided) processes, defined as (30) and (31) 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. Note that the worker's effort does not appear in (31). The probability measure P^0 gives the law of motion assuming zero effort from the worker. The signal processes are defined independently from the worker's effort. Instead of defining the signals themselves, the worker's effort defines the law of motion of the signals.

This brings us to the second subtlety, standard in continuous-time principal-agent models. To avoid circularity problems where actions depend on the process that they define, (11) and (26) 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 fixed once and for all *ex ante*. They are defined for a reference effort level, zero effort, as in (31). The worker's actions do not define the signal process itself. Instead, they define the law of the process: given $A \in \mathcal{A}$, define Z_k^A by $Z_{k,t}^A = Z_{k,t} - \frac{\alpha_k}{\sigma_k^2} \int_0^t A_s ds$. By the Girsanov Theorem, there exists a probability measure P^A such that the joint law of $(\theta, Z_1^A, \dots, Z_K^A)$ under P^A is the same as the joint law of $(\theta, Z_1, \dots, Z_K)$ under P^0 . Given $A \in \mathcal{A}$, the signal S_k satisfies

$$dS_{k,t} = (\alpha_k A_t + \beta_k \theta_t) dt + \sigma_k dZ_{k,t}^A,$$

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

Finally, when we write $\int_{s \leq t} f(s) ds$ for some function f , we adopt the following convention: in the nonstationary setting, $\int_{s \leq t} f(s) ds = \int_0^t f(s) ds$, while in the stationary setting, $\int_{s \leq t} f(s) ds = \int_{-\infty}^t f(s) ds$. The same convention holds for stochastic integrals.

B Formal Statements: Complements

In this appendix, we provide formal statements for claims that are either missing or 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).

Unless mentioned otherwise, the stationary setting is assumed.

Recall that, in the most general form of the model, the rater supplies arbitrary information to the market, and so defines the *information structure* of the market, which is a collection of sigma-algebras $\{\mathcal{M}_t\}_{t \geq 0}$, where \mathcal{M}_t is the market's information at time t . Information may be confidential or public as in Section 5.1, and it may be exclusive or nonexclusive 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, 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 nonexclusive information with public signals S_1, \dots, 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 . This defines public and nonexclusive information structures.

From Section 3, we know that equilibrium behavior is determined by the market's belief about the worker's ability $(\mu_t)_t$, which the market computes based on its information $(\mathcal{M}_t)_t$. As explained, one need not specify the complete information structure—the market belief itself suffices.

For each information structure corresponding to the four cases (confidential/public, exclusive/nonexclusive), we consider the equivalent environment in which the rater provides a single scalar rating which is linear or affine in the market belief. Ratings are public when they are equal to the market belief derived from a public information structure, or a nonconstant affine transformation thereof. Ratings are nonexclusive when they are the market belief derived from a nonexclusive information structure, or a nonconstant affine transformation thereof. Naturally, ratings that are public and/or nonexclusive must satisfy some conditions that ratings under the baseline, confidential exclusive information, need not satisfy. In our statements below, we describe the conditions required for a rating to be a market belief with or without public/nonexclusive constraints.

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.

B.1 Statements of Section 3

We start with the informal claim made in (13) and (14) 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 information structure if, and only if, for all t ,

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

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

In ft.19, we mentioned that the exact information that the worker receives is irrelevant. The next lemma makes this precise. The lemma applies to both stationary and nonstationary settings.

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 23. This follows from the next, more general lemma.

Lemma B.3 Let Y be a linear stationary rating with normalized variance, $\mathbf{Var}[Y_t] = 1$.⁵⁶ The unique equilibrium effort level A is constant and determined by

$$c'(A) = \frac{\gamma^2}{2} \left[\sum_{k=1}^K \alpha_k \int_0^\infty u_k(t) e^{-rt} dt \right] \left[\sum_{k=1}^K \beta_k \int_0^\infty u_k(t) e^{-t} dt \right]. \quad (32)$$

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 (32) in a compact way as

$$c'(A^*) = \left[\sum_{k=1}^K \alpha_k \int_0^\infty u_k(t) e^{-rt} dt \right] \frac{\mathbf{Cov}[Y_t, \theta_t]}{\mathbf{Var}[Y_t]}. \quad (33)$$

Finally, we state below the general version of the Representation Theorem 3.6 for the multisignal framework of Section 5.⁵⁷

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 $s > 0$, $(Y_t, \mathbf{S}_t - \mathbf{S}_{t-s}) - \mathbf{E}[(Y_t, \mathbf{S}_t - \mathbf{S}_{t-s})]$ 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 , $k = 1, \dots, K$, such that, up to an additive constant,

$$Y_t = \sum_{k=1}^K \int_{s \leq t} u_k(t-s) dS_{k,s}. \quad (34)$$

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 $t \geq 0$,

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

⁵⁶The somewhat unwieldy statement of this constraint in terms of $\{u_k\}_k$ is given in (40) below.

⁵⁷The restriction to stationary processes is key to obtaining a linear representation, given the continuum of equations that calls for stronger assumptions. See Jeulin and Yor (1979) for a counterexample otherwise.

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. Stationary is not required, and so to simplify, the proposition assumes the nonstationary setting.

Proposition B.5 *Let Y be a (one-sided) rating process progressively measurable w.r.t. \mathcal{R} . Suppose that the worker employs a deterministic effort strategy, and that under such effort, $(Y_t, X_t - X_{t-s}) - \mathbf{E}[(Y_t, X_t - X_{t-s})]$ is jointly Gaussian in t for $t \geq s$, and Y is such that the following conditions are satisfied:*

1. $\mathbf{Cov}[Y_t, X_\tau | \mathcal{R}_s]$ is differentiable in τ with a derivative that is uniformly Lipschitz continuous in s ;
2. $\mathbf{Cov}[Y_t, X_\tau | \mathcal{R}_s]$ and $\mathbf{Cov}[Y_t, \theta_s | \mathcal{M}_s]$ are deterministic, for all $t > \tau > s$;⁵⁸
3. $\mathbf{E}[Y_t^2] < \infty$ and $\mathbf{E}[|Y_t|] < \infty$.

Then, there exists a function $u_{s,t}$, $s \leq t$, with $s \mapsto u_{s,t}$ integrable and square integrable, such that, for all t ,

$$Y_t = \int_0^t u_{s,t} dX_s, \quad (\text{up to an additive constant function of time}).$$

B.2 Statements of Section 4

In the text and ft.32, 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

$$Y_t = \int_{s \leq t} e^{-\delta(t-s)} [c dX_s + (1-c) dS_s],$$

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

$$Y_t = \int_{t-T}^t [c dX_s + (1-c) dS_s],$$

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.*

⁵⁸That is, nonrandom functions of (s, t, τ) (resp., (s, t)).

As mentioned in ft.30, 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*

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

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

$$m_\alpha = \sum_{k=1}^K \frac{\alpha_k^2}{\sigma_k^2}, \quad m_{\alpha\beta} = \sum_{k=1}^K \frac{\alpha_k \beta_k}{\sigma_k^2}, \quad m_\beta = \sum_{k=1}^K \frac{\beta_k^2}{\sigma_k^2}. \quad (35)$$

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

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

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

$$u_k^c(t) = \frac{\beta_k}{\sigma_k^2} \left(d_k^c \frac{\sqrt{r}}{\lambda} e^{-rt} + e^{-\kappa t} \right), \quad (36)$$

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 (36) 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 $\tau \geq 0$,*

$$\mathbf{Corr}[Y_t, Y_{t+\tau}] = \mathbf{Corr}[\theta_t, \theta_{t+\tau}] (= e^{-\tau}). \quad (37)$$

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

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

Equation (37) captures the constraint that the (belief) rating at any time must convey all relevant information on the worker's ability that is included in the rating disclosed at previous times.

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 NonExclusive Ratings

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

Proposition B.11 (NonExclusive Belief) *Let Y be a stationary linear rating. Then, Y is:*

1. *A belief for a confidential information structure with nonexclusive signals S_1, \dots, S_{K_0} if, and only if, for all t ,*

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

2. *A belief for a public information structure with nonexclusive signals S_1, \dots, S_{K_0} if, and only if, for all t ,*

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

Equivalently, a rating Y is a belief for a confidential/public information structure with nonexclusive signals S_1, \dots, S_{K_0} if, and only if, it is a belief for a confidential/public information structure and, for all $k = 1, \dots, K_0$, all t , and all $\tau \geq 0$,

$$\mathbf{Cov}[S_{k,t}, Y_{t+\tau}] = \mathbf{Cov}[S_{k,t}, \theta_{t+\tau}]. \quad (38)$$

Equation (38) captures the constraint that the (belief) rating at any time must convey all relevant information on the worker's ability that is included in the signals S_1, \dots, S_{K_0} up to that time.

Giving the explicit solution for the optimal rating under nonexclusivity 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 nonexclusive 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, nonexclusive information.

Theorem B.13 *The optimal nonexclusive public rating is unique and given by, for signals $k \leq K_0$,*

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

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 OA.10 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 nonexclusive 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 nonstandard. Hopefully, it may be useful in related contexts.

The proof of each of the four cases considered (confidential/public and exclusive/nonexclusive) 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 nonexclusive. 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 (40) 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

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

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} :

$$\frac{\gamma^2}{2} \left[\sum_{k=1}^K \alpha_k \int_0^\infty u_k(t) e^{-rt} dt \right] \left[\sum_{k=1}^K \beta_k \int_0^\infty u_k(t) e^{-t} dt \right]. \quad (39)$$

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:

$$\sum_{k=1}^K \sigma_k^2 \int_0^\infty u_k(s)^2 ds + \frac{\gamma^2}{2} \int_0^\infty \int_0^\infty U(s)U(t) e^{-|s-t|} ds dt = 1, \quad (40)$$

where $U := \sum_k \beta_k u_k$. The higher dimensionality of the problem is plain in (40). Maximizing (39) subject to (40) 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 (39) and constraint (40) 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 Part VI 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 nonexclusive settings (see Section 5.2), the objective (39) remains the same, but there are additional constraints on the rating. These capture the restriction

that market beliefs are linked to public or nonexclusive ratings. Propositions B.1 and B.9 state these constraints in the exclusive case, and Proposition B.11 does so for the nonexclusive 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.⁵⁹ For example, in the public exclusive setting, the constraint (40) is replaced by

$$\sum_{k=1}^K \sigma_k^2 \int_0^\infty u_k(t) u_k(t + \tau) dt + \frac{\gamma^2}{2} \int_0^\infty \int_0^\infty U(s) U(t) e^{-|s+\tau-t|} ds dt = 1, \quad \forall \tau \geq 0. \quad (41)$$

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 public exclusive setting, we replace (41) by

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

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

Part II: Verification

The calculus of variations determines an essentially unique candidate for the filter \mathbf{u} and thus a unique candidate rating. However, few sufficient conditions are known in the calculus of variations. Most are based on the Hilbert Invariant Integral. However, in the case of (even one-dimensional) integral equations with delayed argument, the method does not apply (Sabbagh (1969)).⁶⁰ Instead, we interpret the rater’s optimization differently, as a principal-agent model. In this auxiliary model,

⁵⁹There 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.

⁶⁰The 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

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 ,

$$\mathbf{E} \left[\int_{s \geq t} e^{-r(s-t)} (\mu_s - c(A_s)) ds \mid \mathcal{R}_t \right]. \quad (42)$$

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

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

$$\mathbf{E} \left[\int_0^\infty \rho e^{-\rho t} (c'(A_t) + \phi \mu_t (\nu_t - \mu_t)) dt \right], \quad (43)$$

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 (42).

To interpret the principal’s objective, it is useful to consider the reward appearing in (43). 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 $\mathbf{E}[\mu_t (\nu_t - \mu_t)] = 0$: 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 $\rho \rightarrow 0$.

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

space, which in our case is not tractable.

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 nonexclusive 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 nonexclusive 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 $\rho \rightarrow 0$. 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, 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 result applies to the nonstationary setting, in which case we recall that $\{\mathcal{R}_t\}_{t \geq 0}$ is simply the filtration generated by the output process X , which is a (one-sided) process starting at time 0.

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 (a Malliavin derivative) has become common in the literature on dynamic contracting following notably the seminal work of Sannikov (2008). 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 A^* be the worker's (deterministic) strategy conjectured by both the rater and the market, which may differ from A^\dagger . 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

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

Similarly, let

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

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

$$\nu_t^* = \int_0^t e^{-\int_s^t (1+m(s')) ds'} m(s) (dX_s - A_s^* ds).$$

Alternatively, we may also write ν_t^* as

$$\nu_t^* = \int_0^t e^{-(t-s)} m(s) dB_s^*.$$

Let $\psi_{t,s} = e^{-(t-s)} m(s)$. An analogous set of equations holds for ν^\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

$$X_t = B_t^* + \int_0^t A_s^* ds + \int_0^t \nu_s^* ds = B_t + \int_0^t A_s^* ds + \int_0^t \int_0^i e^{-(i-j)} m(j) dB_j^* di,$$

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

Next, we let

$$V_\infty = \int_0^\infty e^{-rt} (\mu_t - c(A_t^*)) dt,$$

where $\mu_t = \mathbf{E}^*[\theta_t \mid \mathcal{M}_t]$, and

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

Observe that V is a (bounded) \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

$$V_t = V_0 + \int_0^t \sigma e^{-rs} \xi_s dB_s^*, \quad \text{with } V_0 = - \int_0^\infty e^{-rt} c(A_t^*) dt,$$

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^* :

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

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

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

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

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

and, since

$$V_t = \int_0^t e^{-rs} (\mu_s - c(A_s^*)) ds + e^{-rt} W_t,$$

we have

$$\begin{aligned} d(e^{-r(s-t)} W_s) &= e^{rt} (dV_s - e^{-rs} \mu_s ds + e^{-rs} c(A_s^*) ds) \\ &= e^{-r(s-t)} (\sigma \xi_s dB_s^* - \mu_s ds + c(A_s^*) ds). \end{aligned}$$

Thus,

$$e^{-r\delta} \mathbf{E}^\dagger [W_{t+\delta} \mid \mathcal{R}_t] = W_t + \mathbf{E}^\dagger \left[\int_t^{t+\delta} e^{-r(s-t)} (\sigma \xi_s dB_s^* - \mu_s ds + c(A_s^*) ds) \mid \mathcal{R}_t \right].$$

We also have the equality

$$dX_s = \sigma dB_s^* + (A_s^* + \nu_s^*) ds = \sigma dB_s^\dagger + (A_s^\dagger + \nu_s^\dagger) ds,$$

and hence,

$$\sigma \xi_s dB_s^* = \sigma \xi_s dB_s^\dagger + \xi_s (A_s^\dagger - A_s^*) ds + \xi_s (\nu_s^\dagger - \nu_s^*) ds.$$

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

$$\begin{aligned} &e^{-r\delta} \mathbf{E}^\dagger [W_{t+\delta} \mid \mathcal{R}_t] \\ &= W_t + \mathbf{E}^\dagger \left[\int_t^{t+\delta} e^{-r(s-t)} (\xi_s (A_s^\dagger - A_s^*) + \xi_s (\nu_s^\dagger - \nu_s^*) - \mu_s ds + c(A_s^*)) ds \mid \mathcal{R}_t \right]. \end{aligned}$$

Altogether,

$$W'_t = W_t + \mathbf{E}^\dagger \left[\int_t^{t+\delta} e^{-r(s-t)} (c(A_s^*) - c(A_s^\dagger) + \xi_s(A_s^\dagger - A_s^*) + \xi_s(\nu_s^\dagger - \nu_s^*)) ds \middle| \mathcal{R}_t \right].$$

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

$$\mathbf{E}^\dagger \left[\int_t^{t+\delta} e^{-r(s-t)} (c(A_s^*) - c(A_s^\dagger) + \xi_s(A_s^\dagger - A_s^*) + \xi_s(\nu_s^\dagger - \nu_s^*)) ds \middle| \mathcal{R}_t \right] \leq 0,$$

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

We have:

$$\nu_s^\dagger - \nu_s^* = \int_t^s e^{-\int_j^s (1+m(s')) ds'} m(j) (A_j^\dagger - A_j^*) dj.$$

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

$$\mu_t = \int_0^t \phi_{t,s} dB_s^*.$$

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

$$Y_t^* = \int_0^t |\bar{\phi}_{t,s}| dB_s^*.$$

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

The belief process induced by the confidential rating Y^* is

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

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). Also note that K_t does not matter for the worker's incentives. Hence, it is irrelevant in the sequel.

We have

$$\mathbf{Cov}[Y_t^*, \theta_t] = \mathbf{E}^*[Y_t^* \theta_t] = \mathbf{E}^*[Y_t^* \nu_t^*] = \mathbf{E}^* \left[\int_0^t |\bar{\phi}_{t,s}| \psi_{t,s} ds \right] \geq \mathbf{E}^* \left[\int_0^t \phi_{t,s} \psi_{t,s} ds \right] = \mathbf{E}^*[\mu_t \nu_t^*],$$

because $\psi_{t,s} = e^{-(t-s)m(s)} \geq 0$. We also have

$$\mathbf{Var}[Y_t^*] = \mathbf{E}^*[(Y_t^*)^2] = \int_0^t |\bar{\phi}_{t,s}|^2 ds = \int_0^t \mathbf{E}^*[\phi_{t,s}^2] ds \leq \int_0^t \mathbf{E}^*[\phi_{t,s}^2] ds = \mathbf{E}^*[\mu_t^2]. \quad (44)$$

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

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

Hence,

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

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

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

Let ξ^* be the analogue, for belief μ^* , of ξ as defined in Step 1 for belief μ . Recall that, for almost all t , $c'(A_t^*) = \xi_t(1 + m(t))$, and similarly, $c'(A_t^*) = \xi_t^*(1 + m(t))$, so $c'(A_t^*)/c'(A_t^*) = \xi_t^*/\xi_t$.

The process ξ can be written explicitly as,

$$\xi_t = \frac{1}{\sigma} \int_t^\infty e^{-r(s-t)} \phi_{s,t} ds$$

and similarly

$$\xi_t^* = \frac{1}{\sigma} \int_t^\infty e^{-r(s-t)} \frac{\mathbf{Cov}[\theta_t, Y_t^*]}{\mathbf{Var}[Y_t^*]} |\bar{\phi}_{s,t}| ds.$$

For almost every t , ξ_t must be deterministic, so

$$\frac{1}{\sigma} \int_t^\infty e^{-r(s-t)} \phi_{s,t} ds = \frac{1}{\sigma} \int_t^\infty e^{-r(s-t)} \bar{\phi}_{s,t} ds,$$

and as $|\bar{\phi}_{s,t}| \geq \max\{0, \bar{\phi}_{s,t}\}$ and $\mathbf{Cov}[\theta_t, Y_t^*]/\mathbf{Var}[Y_t^*] \geq 1$, we get $\xi_t^* \geq \xi_t$ implying $A_t^* \geq A_t^*$, for almost all t .

This proof has two implications. First, to reach the maximum possible effort almost everywhere, a linear rating (such as the market belief of any rating, by its martingale representation) must have nonrandom weights almost everywhere. If the weights of a rating are random, then the inequality in Equation (44) becomes strict, and using the transformed rating Y^* yields strictly high effort levels for some positive measure of times. Second, adding external noise to any rating always has the effect of depressing the worker's equilibrium effort. Although such a rating is outside of the model assumed in this proof, it is straightforward to carry the above computations when using rating $Y^* = \mu_t + \tilde{B}_t$ instead, where \tilde{B} is an independent Brownian motion. The effort that result is

depressed because the noisy rating has an increased variance while the covariance of the rating with the worker's ability remains the same.

D.2 Proof of Theorem 3.6

The proof proceeds in two parts. We start with a candidate guess for the function u , whose derivation is explained in Section OA.4 of the Online Appendix, which also includes the proof for the multisignal version of the theorem (Theorem B.4).⁶¹ 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: Let

$$f(s) := \mathbf{Cov}[Y_t, X_{t-s}], \text{ and } \bar{f}(s) := \frac{1}{\sigma^2} f(s).$$

$$C_1^\infty := \frac{\kappa^2 - 1}{2\kappa} \int_0^\infty \bar{f}'(j) e^{-\kappa j} dj, \quad (45)$$

and

$$C_2^\infty := \frac{(\kappa - 1)^2}{2\kappa} \int_0^\infty \bar{f}'(j) e^{-\kappa j} dj. \quad (46)$$

A guess for the candidate u is:

$$u(\tau) := C_1^\infty e^{\kappa\tau} + C_2^\infty e^{-\kappa\tau} - \bar{f}'(\tau) - \frac{\kappa^2 - 1}{\kappa} \int_0^\tau \sinh(\kappa(\tau - s)) \bar{f}'(s) ds, \quad (47)$$

or equivalently,

$$u(\tau) := \frac{\kappa^2 - 1}{\kappa} \left(\frac{\sinh \kappa\tau + \kappa \cosh \kappa\tau}{1 + \kappa} \int_0^\infty e^{-\kappa s} d\bar{f}(s) - \int_0^\tau \sinh \kappa(t - s) d\bar{f}(s) \right) - \bar{f}'(\tau). \quad (48)$$

Proof of Integrability. We show that every u_k defined by Equation (48) is integrable and square integrable. To do so, we have to show that

$$(\sinh \kappa t + \kappa \cosh \kappa t) \int_0^\infty e^{-\kappa s} h(s) ds - (1 + \kappa) \int_0^t \sinh \kappa(t - s) h(s) ds \quad (49)$$

is integrable and square integrable whenever h and h^2 are. We note that (49) is linear in h , so that it suffices to show that its positive and negative parts are integrable. Hence, without loss, we assume

⁶¹Determining 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 nontrivial because it prevents the use of the Kalman-Bucy filter and involves finding a continuum of terms of the form $\mathbf{Var}[Y_t | \mathcal{R}_{t-s}]$ 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.

that $h \geq 0$. After rearranging the terms, (49) is equal to

$$\frac{1}{2}(\kappa + 1) \left(e^{\kappa t} \int_t^\infty e^{-\kappa s} h(s) ds + e^{-\kappa t} \int_0^t e^{\kappa s} h(s) ds \right) + \frac{1}{2}(\kappa - 1) e^{-\kappa t} \int_0^\infty e^{-\kappa s} h(s) ds. \quad (50)$$

Thus, (49) is nonnegative, and showing the integrability of (49) reduces to showing that the integral of (49) converges on $[0, +\infty)$. It is readily verified by differentiation that (49) is the derivative of

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

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

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

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

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

converges, which is immediate from the integrability of h . Thus, (49) is integrable. Next, to show that (50) is square integrable, we show that

$$e^{\kappa t} \int_t^\infty e^{-\kappa s} h(s) ds \quad (51)$$

is square integrable. As square-integrable functions are closed under additivity, and h is integrable, (51) is the only term of (50) for which square integrability is nontrivial. By the Cauchy-Schwarz inequality,

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

Thus,

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

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 defines valid coefficients for the rating. Let u be defined by (48), or, equivalently, by (47). Let

$$\tilde{Y}_t := \mathbf{E}^*[Y_t] + \int_{s \leq t} u(t-s)(dX_s - A_s^* ds).$$

If we have $\mathbf{Cov}[Y_t - \tilde{Y}_t, X_{t-\tau}] = 0$ for all τ , then Y_t and $X_{t-\tau}$ are independent for all τ . As $Y_t - \tilde{Y}_t$ is measurable with respect to the information generated by the outputs $X_{t-\tau}$, $\tau \geq 0$, it implies that $\mathbf{Var}[Y_t - \tilde{Y}_t] = 0$ and thus $Y_t = \tilde{Y}_t$. In the remainder, we show that $\mathbf{Cov}[Y_t - \tilde{Y}_t, X_{t-\tau}] = 0$ for all $\tau \geq 0$. Let $g(\tau) = \mathbf{Cov}[\tilde{Y}_t, X_{t-\tau}]$. We have:

$$g(\tau) = \int_0^\infty u(s) \mathbf{Cov}[dX_{t-s}, X_{t-\tau}] = \sigma^2 \int_\tau^\infty u(s) ds + \frac{\gamma^2}{2} \int_0^\infty \int_\tau^\infty u(s) e^{-|s-j|} dj ds,$$

and so

$$g'(\tau) = -\sigma^2 u(\tau) - \frac{\gamma^2}{2} \int_0^\infty u(s) e^{-|\tau-s|} ds.$$

Replacing u by its definition in (47),

$$\begin{aligned} g'(\tau) &= f'(\tau) - C_1 \frac{\gamma^2}{\kappa^2 - 1} e^{\kappa\tau} - C_2 \frac{\gamma^2}{\kappa^2 - 1} e^{-\kappa\tau} + \frac{\gamma^2}{\kappa} \int_0^\tau \sinh(\kappa(\tau-s)) \bar{f}'(s) ds \\ &\quad - \frac{\gamma^2}{2} \int_0^\infty u(s) e^{-|\tau-s|} ds. \end{aligned} \tag{52}$$

From (47), we have

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

It holds that

$$\int_0^\infty u(s) e^{-|\tau-s|} ds = \lim_{B \rightarrow \infty} \int_0^B u(s) e^{-|\tau-s|} ds.$$

Thus,

$$\begin{aligned} \int_0^B u(s) e^{-|\tau-s|} ds &= C_1^\infty \int_0^B e^{\kappa s} e^{-|\tau-s|} ds + C_2^\infty \int_0^B e^{-\kappa s} e^{-|\tau-s|} ds - \int_0^B \bar{f}'(s) e^{-|\tau-s|} ds \\ &\quad - \frac{\kappa^2 - 1}{\kappa} \int_0^B \int_0^s \sinh(\kappa(s-j)) \bar{f}'(j) e^{-|\tau-s|} dj ds. \end{aligned}$$

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

$$\begin{aligned} \int_0^B \int_0^s \sinh(\kappa(s-j)) \bar{f}'(j) e^{-|\tau-s|} dj ds &= -\frac{\kappa}{\kappa^2-1} \int_0^B \bar{f}'(j) e^{-|\tau-j|} dj \\ &+ \frac{e^{-B+\tau}}{\kappa^2-1} \int_0^B \bar{f}'(j) [\kappa \cosh(\kappa(B-j)) + \sinh(\kappa(B-j))] dj - \frac{2}{\kappa^2-1} \int_0^\tau \sinh(\kappa(\tau-j)) \bar{f}'(j) dj. \end{aligned}$$

Using the expressions for C_1^∞ and C_2^∞ given by (45) and (46), we get that

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

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

$$\begin{aligned} \int_0^\infty u(s) e^{-|\tau-s|} ds &= \frac{2}{\kappa} \int_0^\tau \sinh(\kappa(\tau-j)) \bar{f}'(j) dj \\ &+ C_1^\infty \left[\frac{e^{\kappa\tau}}{\kappa+1} - \frac{e^{\kappa\tau}}{\kappa-1} \right] + C_2^\infty \left[\frac{e^{-\kappa\tau}}{\kappa+1} - \frac{e^{-\kappa\tau}}{\kappa-1} \right]. \end{aligned} \quad (53)$$

Plugging the expression of (53) in (52) yields

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

So $g' = f'$. As $f(0) = g(0) = 0$, 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 distribution over ratings and outputs.