

*Who Should Pay for Credit Ratings and How?**

Anil K Kashyap[†] and Natalia Kovrijnykh[‡]

October 2012

Abstract

This paper analyzes a model where investors use a credit rating to decide whether to finance a firm. The rating quality depends on the credit rating agency's (CRA) effort, which is unobservable. We analyze optimal compensation contracts for the CRA that differ depending on whether the firm, investors or a social planner orders the rating. We find that rating errors are larger when the firm orders it than when investors do. However, investors ask for ratings inefficiently often. Which arrangement leads to a higher social surplus depends on the agents' prior beliefs about the project quality. We also show that competition among CRAs causes them to reduce their fees, put in less effort, and thus leads to less accurate ratings. Rating quality also tends to be lower for new securities. Finally, we find that optimal contracts that provide incentives for both initial ratings and their subsequent revisions can lead the CRA to be slow to acknowledge mistakes.

Keywords: Rating Agencies, Optimal Contracts, Moral Hazard, Information Acquisition

JEL Codes: C73, D82, D83, D86, G24.

*We have benefited from discussions with Hector Chade, Simon Gilchrist, Robert Lucas, Marcus Opp, Chris Phelan, Robert Shimer, Nancy Stokey, and Joel Watson. We are also grateful for comments by seminar participants at Arizona State University, University of Chicago, and University of Wisconsin—Madison. Kashyap thanks the National Science Foundation and the Initiative on Global Markets at Chicago Booth for research support. For information on Kashyap's outside compensated activities see <http://faculty.chicagobooth.edu/anil.kashyap/>.

[†]Booth School of Business, University of Chicago, and NBER. Email: anil.kashyap@chicagobooth.edu

[‡]Department of Economics, Arizona State University. Email: natalia.kovrijnykh@asu.edu

1 Introduction

Virtually every government inquiry into the 2008 and 2009 financial crisis has assigned some blame to credit rating agencies. For example, the Financial Crisis Inquiry Commission (2011, p. xxv) concludes that “this crisis could not have happened without the rating agencies”. While the United States Senate Permanent Subcommittee on Investigations (2011, p. 6) states that “inaccurate AAA credit ratings introduced risk into the U.S. financial system and constituted a key cause of the financial crisis”. But the details of the indictments differ across the reports. For instance, the Senate report points to inadequate staffing as a critical factor, while the Financial Crisis Inquiry Commission highlights the business model that had firms seeking to issue securities pay for ratings as a major contributor.¹ In this paper we explore the role that these and other factors might play creating inaccurate ratings.

In the environment we study, a firm seeks a credit rating to help it attract investors. The ratings agency creates value by generating information about the investment project that can lead to more efficient financing decisions. But a CRA must exert effort to learn about the quality of a project that a firm seeks to finance. The CRA’s effort is not observable, so the compensation scheme must provide incentives for the CRA to exert effort. This very simple framework allows us to compare the CRA’s effort and the total surplus to all actors in the model depending who orders the rating — a social planner, the firm, or potential investors.

This framework makes it possible to address directly the claims made in the government reports. In particular, we can ask: how do you compensate CRAs to avoid shirking? Does the issuer-pays model generate more shirking than when the investors pay for ratings? In addition, in natural extensions of the basic model we can see whether a battle for market share would be expected to reduce ratings quality, or whether different types of securities create different incentives to shirk.

¹The United States Senate Permanent Subcommittee on Investigations (2011) reported that “factors responsible for the inaccurate ratings include rating models that failed to include relevant mortgage performance data, unclear and subjective criteria used to produce ratings, a failure to apply updated rating models to existing rated transactions, and a failure to provide adequate staffing to perform rating and surveillance services, despite record revenues.” Financial Crisis Inquiry Commission (2011) concluded that “the business model under which firms issuing securities paid for their ratings seriously undermined the quality and integrity of those ratings; the rating agencies placed market share and profit considerations above the quality and integrity of their ratings.”

The framework also explains five facts about the ratings business in a unified fashion. The first fact is that rating mistakes are in part due to insufficient effort by rating agencies. The second is that outcomes and accuracy of ratings do differ depending on which party pays for a rating. Third, increases in competition between rating agencies are accompanied by a reduction in the accuracy of ratings. Fourth, ratings mistakes are more common for newer securities with shorter histories that can be studied than for more established types of securities. Finally, revisions to ratings are slow.

We begin to account for these facts by studying a one-period environment where a single firm is seeking funding for a project from investors. The project's quality is unknown and a single rating agency can be hired to generate information about the project. We characterize the optimal compensation arrangement for the CRA allowing different agents to order a rating. As is often the case in this kind of problem, the need to induce effort argues for giving the surplus from the investment project to the ratings agency, so that the higher CRA's profits, the higher the effort it exerts.

We find that under the issuer-pays model, the rating is acquired less often and is less informative (i.e., the CRA expends less effort) than in the investor-pays model (or in the second best). However, the total surplus in the issuer-pays may be higher or lower than in the investor-pays model, depending on the agents' prior beliefs about the quality of the project. The ambiguity about the total surplus comes because even though investors induce the CRA to devote more effort on ratings, they will ask for ratings even when a social planner would not. So the extra accuracy achieved by having investors pay is potentially dissipated by an excessive reliance on ratings.

We then extend the basic setup in three ways. The first extension explores the implications of allowing rating agencies to compete for business. An immediate implication of competition is a tendency to reduce fees in order to win business. But with lower fees comes lower effort on project evaluation. Hence, this framework predicts that competition will lead to less accurate ratings.

A second extension considers the accuracy of ratings on different types of securities. We suppose that some types of investment projects are inherently more difficult for the CRA to evaluate — presumably because they have a short track record that makes comparisons difficult. In this case it is inevitable that the ratings will deteriorate.

Finally, we allow for a second period in the model and posit that investment is needed in each of the two periods, so that there is a role for ratings in both periods. The need to elicit effort in both periods poses a problem. An efficient compensation arrangement will

reward the CRA for getting a correct rating that stays the same in both periods — this is the best way to create powerful incentives in the first period. But if a mistake is made in the initial rating, the optimal way to create incentives for the last period is to recognize the mistake and change the rating. Balancing these considerations can lead the CRA to be slow to acknowledge mistakes.

While we find that this basic framework is very powerful in that it explains the five aforementioned facts using relatively few assumptions, our approach does come with several limitations. For instance, we have fully rational investors and only one rating agency, so we cannot address debates related to rating shopping — a common criticism of the issuer-pays model.² Likewise, we also assume that the firm has the same knowledge about the quality project *ex ante* as everyone else. Without this assumption, the analysis becomes much more complicated, since in addition to the moral hazard problem on the side of the CRA there is an adverse selection problem on the side of the firm. We do offer some cursory thoughts on this problem in our conclusions.

The remainder of the paper is organized as follows. The next section documents the empirical regularities that motivate our analysis and compares our model to others in the literature. Section 3 introduces the baseline model. Section 4 presents our main propositions about CRA compensation, as well as comparison between the issuer-pays and investor-pays models. Section 5 covers the three extensions just described. Section 6 concludes.

2 Motivating Facts and Literature Review

Given the intense interest in the causes of the financial crisis and the role that official accounts of the crisis ascribe to the ratings agencies, it is not surprising that there has an explosion of research on credit rating agencies. White (2010) offers a concise description of the ratings industry and recounts its role in the crisis. To understand our contribution, we find it helpful to separate the recent literature into three sub-areas.

The first consists of the empirical studies that seek to document mistakes or perverse ratings outcomes. There are so many of these papers that we cannot cover them all, but it is helpful to note that there are five facts that our analysis takes as given. So we will point to specific contributions that document these particular facts.

²See the literature review below for discussion of papers that do generate rating shopping. Notice, however, that even without rating shopping we were able to identify problems with the issuer-pays model.

First, the question of who pays for a rating does seem to matter. Finding recent direct evidence is hard because the issuer-pays model for most securities is so pervasive. But, Jiang, Stanford, and Xie (2012) address this question by analyzing data from the 1970s when Moody's and S&P were using different compensation models. In particular, from 1971 until June 1974 S&P was charging investors for ratings, while Moody's was charging issuers. During this period the Moody's ratings systematically exceeded those of S&P. S&P adopted the issuer-pays model in June 1974 and from that point forward over the next three years their ratings matched Moody's.

Second, as documented by Mason and Rosner (2007), most of the rating mistakes occurred for structured products that primarily related to asset-backed securities. As Pagano and Volpin (2010) note, the volume of these new securities increased tenfold between 2001 and 2010. As Mason and Rosner emphasize, the mistakes that happened for these new products were not found for corporate bonds where the CRAs had much more experience. In addition, Morgan (2002) argues that banks (and insurance companies) are inherently more opaque than other firms, and this opacity explains his finding that Moody's and S&P differ more in their ratings for these intermediaries than for non-banks.

Third, some of the mistakes in the structured products seem to be due to insufficient monitoring and effort on the part of the analysts. For example, Owusu-Ansah (2012) shows that downgrades by Moody's tracked movements in aggregate Case-Shiller home price indices much more than any private information that the agencies had about specific deals.

Interestingly, the Dodd-Frank Act in the U.S. also presumes that shirking was a problem during the crisis and takes several steps to try to correct it. First, section 936 of the Act requires the Securities and Exchanges Commission to take steps to guarantee that any person employed by a nationally recognized statistical rating organization (1) meets standards of training, experience, and competence necessary to produce accurate ratings for the categories of issuers whose securities the person rates; and (2) is tested for knowledge of the credit rating process. The law also requires the agencies to identify and then notify the public and other users of ratings which five assumptions would have the largest impact on their ratings in the event that they were incorrect.

Fourth, revisions to ratings are typically slow to occur. This issue attracted considerable attention early in the last decade when the rating agencies were slow to identify problems at Worldcom and Enron ahead of their bankruptcies. But, Covitz and Harrison (2003) show that 75% of the price adjustment of a typical corporate bond in the wake of a downgrade

occurs prior to the announcement of the downgrade. So these delays are pervasive.

Finally, it appears that competition among rating agencies reduces the accuracy of ratings. Very direct evidence on this comes from Becker and Milbourn (2011) who study how the rise in market share by Fitch influenced ratings by Moody's and S&P (who had historically dominated the industry). Prior to its merger with IBCA in 1997, Fitch had a very low market share in terms of ratings. Thanks to that merger, and several subsequent acquisitions over the next five years, Fitch substantially raised its market share, so that by 2007 it was rating around 1/4 of all the bonds in a typically industry. Becker and Milbourn exploit the cross-industry differences in Fitch's penetration to study competitive effects. They find an unusual pattern. Any given individual bond is more likely to be rated by Fitch when the ratings from the other two big firms are relatively low. Yet, in the sectors where Fitch issues more ratings, the overall ratings for the sector tend to be higher.

This pattern is not easily explained by the usual kind of catering that the rating agencies have been accused of. If Fitch were merely inflating its ratings to gain business with the poorly performing firms, the Fitch intensive sectors would be ones with more ratings for these under-performing firms and hence lower overall ratings. This general increase in ratings suggests instead a broader deterioration in the quality of the ratings, which would be expected if Fitch's competitors saw their rents declining; consistent with this view, the forecasting power of the ratings for defaults also decline.

Our paper is also related to the many theoretical papers on rating agencies that have been proposed to explain these and other facts.³ However, we believe our paper is the only one that simultaneously accounts for the five facts described above.

Two papers that are perhaps closest to ours are Opp, Opp, and Harris (2012) and Bolton, Freixas, and Shapiro (2012). Opp, Opp, and Harris explain rating inflation by building a model where ratings not only provide information to investors, but are also used for regulatory purposes. They find that introducing rating-contingent regulation leads the rating agency to rate more firms highly, although it may increase or decrease rating informativeness. Similar to us, they have a rational expectations model where the CRA's effort affects rating precision. But unlike us, they assume that the CRA can commit to exert effort (or, equivalently, that effort is observable), and they do not study optimal contracts.

³ While not applied to rating agencies, there are a number of theoretical papers on delegated information acquisition, see, for example, Chade and Kovrijnykh (2012), Inderst and Ottaviani (2009, 2011) and Gromb and Martimort (2007). Our paper is also broadly related to the literature on media biases — see, e.g., Mullainathan and Shleifer (2005) and references therein.

Bolton, Freixas, and Shapiro (2012) study a model where a CRA receives a signal about a firm’s quality, and can misreport it (although investors learn about a lie ex post). Some investors are naive which creates incentives for the CRA — which is paid by the issuer — to inflate ratings. The authors show that CRAs are more likely to inflate (misreport) ratings in booms, when there are more naive investors, and/or when the risks of failure which could damage CRA reputation are lower. In their model, both the rating precision and reputation costs are exogenous. In contrast, in our model the rating precision is chosen by the CRA; also, our optimal contract with performance-contingent fees can be interpreted as the outcome of a system in which reputation is endogenous. As in our model, the authors predict that competition among CRAs may reduce market efficiency, but for a very different reason than we do: the issuer has more opportunities to shop for ratings and to take advantage of naive investors by only purchasing the best ratings. In contrast, we assume rational expectations, and predict that larger rating errors occur because of more shirking by CRAs.

Skreta and Veldkamp (2009) analyze a model where the naïveté of investors gives issuers incentives to shop for ratings by approaching several rating agencies and publishing only favorable ratings.⁴ They show that a systematic bias in disclosed ratings is more likely to occur for more complex securities — a finding that resembles our result that rating errors are larger for new securities. Similar to our findings, in their model, competition also worsens the problem. They also show that switching to the investor-pays model alleviates the bias, but as in our set up free rider problems can then potentially eliminate the ratings market completely.

Similar to us, Faure-Grimaud, Peyrache, and Quesada (2009) study optimal contracts between a rating agency and a firm, but their focus is on providing incentives to the firm to reveal its information, while we focus on providing incentives to the CRA to exert effort.⁵ Goel and Thakor (2011) have a model where the CRA’s effort is unobservable, but they do not analyze optimal contracts; instead, they are interested in the impact of legal liability for “misrating” on the CRA’s behavior.

Finally, our analysis is also relevant for the many policy-oriented papers that discuss potential reforms of the agencies. Medvedev and Fennell (2011) provide an excellent summary of these issues. Their survey is also representative of most of the papers on this topic

⁴See also Sangiorgi and Spatt (2012) who generate rating shopping in a model with rational investors.

⁵More generally, the structure of our optimal contracts can be endogenously embodying reputational effects. Other papers that model reputational concerns of rating agencies include, for example, Bar-Isaac and Shapiro (2010), Fulghieri, Strobl, and Xia (2011), and Mathis, McAndrews, and Rochet (2009).

in that it identifies the intuitive conflicts of interest that arise from the issuer-pays funding model, and compares them to the alternatives problems that arise under other schemes (such as investor-pays, or having a government agency issue ratings). But all of these analyses are partly limited by the lack of microeconomic foundations underlying the payment models being contrasted. By deriving the optimal compensation contracts we believe we help clarify these kinds of discussions.

3 The Model

We consider a one-period model with one firm, a number ($n \geq 2$) of investors, and one credit rating agency.⁶ All agents are risk neutral and maximize expected profits.

The firm (the issuer of a security) is endowed with a project that requires one unit of investment (in terms of the consumption good) and generates the end-of-period return, which equals y units of the consumption good in the event of success and 0 in the event of failure. The likelihood of success depends on the quality of the project, q .

The quality of the project can be good or bad, $q \in \{g, b\}$, and is unobservable to everyone.⁷ A project of quality q succeeds with probability p_q , where $0 < p_b < p_g < 1$. We assume that $-1 + p_b y < 0 < -1 + p_g y$, so that it is profitable to finance a high-quality project but not a low one. The prior belief that the project is high quality is denoted by γ , where $0 \leq \gamma \leq 1$.

The CRA can acquire information about the quality of the project. It observes a signal $\theta \in \{h, \ell\}$ that is correlated with the project's quality. How informative the signal is about a project's quality depends on the level of effort $e \geq 0$ that the CRA exerts. Specifically,

$$\Pr\{\theta = h|q = g, e\} = \Pr\{\theta = \ell|q = b, e\} = \frac{1}{2} + e,$$

where e is restricted to be between 0 and $1/2$. Note that if effort is zero, the conditional distribution of the signal is the same regardless of the project's quality, and therefore the signal is uninformative. Conditional on the project being of a certain quality, the probability of observing a signal consistent with that quality is increasing in the agent's effort. So higher effort makes the signal more informative in Blackwell's sense.⁸

⁶In Section 5.1, we partially analyze the model with several rating agencies.

⁷We discuss what happens if the issuer has private information about its type in the conclusions.

⁸See Blackwell and Girshick (1954), chapter 12.

Exerting effort is costly for the CRA, where $\psi(e)$ denotes the cost of effort e , in units of the consumption good. The function ψ satisfies $\psi(0) = 0$, $\psi'(e) > 0$, $\psi''(e) > 0$, $\psi'''(e) > 0$ for all $e > 0$, and $\lim_{e \rightarrow 1/2} \psi(e) = +\infty$. The assumptions on the second and third derivatives of ψ guarantee that the CRA's and planner's problems, respectively, are strictly concave in effort. We also assume that $\psi'(0) = 0$ and $\psi''(0) = 0$, which guarantee an interior solution for effort in the CRA's and planner's problems, respectively.

To keep the analysis simple, we will assume that the CRA cannot lie about a signal realization so the rating it announces will be the same as the signal, although we will dispose of this assumption in Section 5.2. In addition, we assume that the CRA is protected by limited liability, so that all payments that it receives must be non-negative.

The firm has no internal funds, and therefore needs investors to finance the project.⁹ Investors behave competitively and are deep-pocketed so that there is never a shortage of funds.

We will consider three scenarios depending on who decides whether a rating is ordered — the social planner, the issuer, or each of the investors. Let X refer to the identity of the player ordering a rating. The timing of events, as shown in Figure 1, is as follows.

At the beginning of each period, the CRA posts a rating fee schedule (the fees to be paid at the end of the period, conditional on the history). Each investor offers a contract (financing terms) to the firm that specifies the amount of investment conditional on a rating or the absence of one, and history-contingent repayments to be made by the firm. The firm either accepts or rejects.¹⁰ X decides whether to ask for a rating, and chooses whether to reveal to the public that a rating has been ordered.¹¹ If a rating is ordered, the CRA exerts effort and announces the rating to X , who then decides whether it should be published (and hence made known to other agents). Investment proceeds according to the contract. If the project is financed, its success or failure is observed. The firm pays investors as specified in the contract, and the CRA collects its fees.¹²

We are interested in analyzing Pareto efficient equilibria in this environment. We will

⁹We make this assumption for expositional convenience. Our results would not change if the firm had initial wealth which is strictly smaller than one — the amount of funds needed to finance the project.

¹⁰We assume that if the firm is indifferent between investors' offers, it either randomizes with equal probabilities over which offer to accept, or, equivalently, obtains an equal amount of funds from each investor.

¹¹We make this assumption only for concreteness. Our results would not change if X could not credibly announce the fact that a rating was ordered, or if everyone automatically observed whether a rating was ordered (but did not necessarily find out what rating was given).

¹²We assume that X can commit to paying the fees due to the CRA, and that the firm can commit to paying investors.

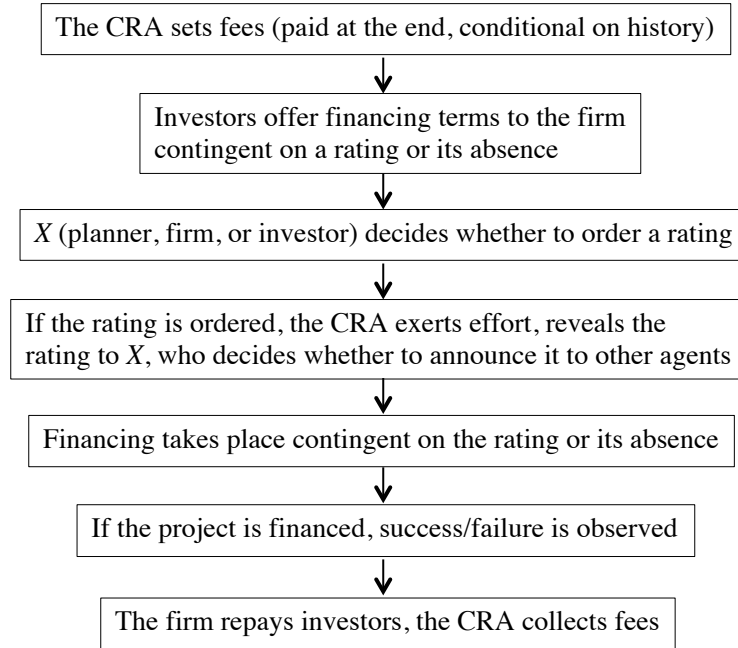


Figure 1: Timing.

compare effort and total surplus depending on who orders a rating. The rationale for considering total surplus comes from thinking about a hypothetical consumer who owns both the firm and CRA, in which case it would be natural for the social planner to maximize the consumer's utility. In our static environment, we will not always be able to Pareto rank equilibria depending on who orders the rating. However, it can be shown that constraints that lead to a lower total surplus in the static model, lead to Pareto dominance in a repeated infinite horizon version of the model.

4 Analysis and Results

It will be convenient to define the following objects. First, let π_1 denote the ex-ante probability of success (before observing a rating), so $\pi_1 = p_g\gamma + p_b(1 - \gamma)$. Next, let $\pi_h(e)$ denote the probability of observing a high rating given effort e , that is, $\pi_h(e) = (1/2 + e)\gamma + (1/2 - e)(1 - \gamma)$. The probability of observing the low rating given effort e is then $\pi_\ell(e) = 1 - \pi_h(e)$. Also, let $\pi_{h1}(e)$ and π_{h0} denote the probabilities of observing a high rating followed by the project's success/failure given effort e : $\pi_{h1}(e) = p_g(1/2 + e)\gamma + p_b(1/2 - e)(1 - \gamma)$ and $\pi_{h0}(e) = (1 - p_g)(1/2 + e)\gamma + (1 - p_b)(1/2 - e)(1 - \gamma)$.

Since producing rating is costly, it cannot be optimal to pay that cost if the information is not used. This implies that if the CRA exerts positive effort, then the project must be financed after the high rating and will not be financed after the low rating.

As a useful benchmark, we consider the first-best case, where the CRA's effort is observable, and the social planner decides whether to order a rating.¹³ There are three cases to consider: (i) do not acquire a rating and do not finance the project, (ii) do not acquire a rating and finance the project, and (iii) acquire a rating and finance the project only if the rating is high. Combining the three options, the total surplus in the first-best case is $S^{FB} = \max\{0, -1 + \pi_1 y, \max_e -\psi(e) - \pi_h(e) + \pi_{h1}(e)y\}$. Define $e^* = \arg \max_e -\psi(e) - \pi_h(e) + \pi_{h1}(e)y$, which is the optimal level of effort under the third alternative. Denote the first-best effort by e^{FB} . Notice that $e^{FB} = 0$ in cases (i) and (ii), and $e^{FB} = e^* > 0$ in case (iii). The following lemma shows which of the three cases occurs depending on the prior γ .

Lemma 1 *There exist thresholds $\underline{\gamma}$ and $\bar{\gamma}$ satisfying $0 < \underline{\gamma} < \bar{\gamma} < 1$, such that*

- (i) $e^{FB} = 0$ for $\gamma \in [0, \underline{\gamma}]$, and the project is never financed;
- (ii) $e^{FB} = 0$ for $\gamma \in [\bar{\gamma}, 1]$, and the project is always financed;
- (iii) $e^{FB} > 0$ for $\gamma \in (\underline{\gamma}, \bar{\gamma})$, and the project is only financed after the high rating.

The intuition behind this result is quite simple. If the prior belief about the project quality is close to zero or one, it does not pay off to acquire additional information about the quality of the project.

The more interesting case is when the CRA's effort is unobservable, and payments are subject to limited liability. The CRA will now choose its effort privately, given the fees it expects to receive at the end of the period.

We will first characterize the (constrained) Pareto frontier in this setup. Depending on which player orders the rating, we will consider an equilibrium where the total surplus is maximized. Each of the equilibria will lie at a different point on this Pareto frontier.¹⁴

In order to construct the Pareto frontier, we need to analyze the optimal contract (fee structure) that provides the CRA with incentives to exert effort. The best way to provide incentives is to pay fees contingent on possible outcomes. If the CRA is asked for a rating, there are three possibilities: the rating is high and the project succeeds, the rating is high

¹³In fact, it is easy to check that when effort is observable, the total surplus is the same regardless of who orders a rating.

¹⁴In fact, in one important case that we explain below, the equilibrium will lie inside the Pareto frontier.

and the project fails, and the rating is low (in which case the project is not financed). Let f_{h1} , f_{h0} , and f_ℓ denote the fees that the CRA receives in each scenario.^{15,16}

On the Pareto frontier, the value to one party is maximized subject to delivering at least certain values to other parties. Investors behave competitively and thus always earn zero profits. Therefore, we can maximize the value to the firm subject to delivering at least a certain value v to the CRA.

As before, there are three options available — do not acquire a rating and do not finance the project, do not acquire a rating and finance the project, and acquire a rating and finance the project only if the rating is high. Let $u(v)$ denote the value to the firm under the third alternative, given that the value to the CRA is at least v . Since investors earn zero profits, the firm extracts all the surplus generated in production, net of the expected fees paid to the CRA. Then the Pareto frontier can be written as $\max\{0 - v, -1 + \pi_1 y - v, u(v)\}$, where

$$u(v) = \max_{e, f_{h1}, f_{h0}, f_\ell} -\pi_h(e) + \pi_{h1}(e)y - \pi_{h1}(e)f_{h1} - \pi_{h0}(e)f_{h0} - \pi_\ell(e)f_\ell \quad (1)$$

$$\text{s.t. } -\psi(e) + \pi_{h1}(e)f_{h1} + \pi_{h0}(e)f_{h0} + \pi_\ell(e)f_\ell \geq v, \quad (2)$$

$$\psi'(e) = \pi'_{h1}(e)f_{h1} + \pi'_{h0}(e)f_{h0} + \pi'_\ell(e)f_\ell, \quad (3)$$

$$e \geq 0, f_{h1} \geq 0, f_{h0} \geq 0, f_\ell \geq 0. \quad (4)$$

Constraint (2) ensures that the CRA's profits are at least v . Constraint (3) is the first-order condition of the CRA's problem, which is obtained by maximizing the left-hand side of (2) with respect to e . Constraints (4) reflect limited liability and nonnegativity of effort.

The following proposition describes the optimal fee structure for the CRA, contingent on the outcomes.

Proposition 1 (Optimal Compensation Scheme) *Suppose the project is financed only after the high rating. Define the cutoff value $\hat{\gamma} = 1/(1 + \sqrt{p_g/p_b})$.*

(i) *If $\gamma \geq \hat{\gamma}$, then it is optimal to set $f_{h1} > 0$ and $f_\ell = f_{h0} = 0$.*

¹⁵One might argue that in reality CRAs are mostly compensated upfront. In the static model, an upfront fee will never provide the CRA with incentives to exert effort — the CRA will take the money and shirk. In a dynamic model, an upfront payment that depends on the past outcomes would be possible. In this case, the subsequent rewards for having a reputation for accuracy would create incentives to exert effort. The fee structure in our static model can then be viewed as a reduced form of such a reputation mechanism.

¹⁶Notice that when X is the firm and acquiring a rating is profitable, investors would be willing to lend to the firm for the rating purchase. Thus the firm is able to pay f_ℓ or f_{h0} , even though the project yields no return after these histories. In particular, if $f_\ell > 0$, investors lend f_ℓ upfront, plus one unit of investment if the rating is high. The firm pays f_ℓ back to investors if the rating is high (and pays it to the CRA otherwise), plus the return specified in the contract if the project succeeds.

(ii) If $\gamma \leq \hat{\gamma}$, then it is optimal to set $f_\ell > 0$ and $f_{h1} = f_{h0} = 0$.

The proposition states that there is a threshold level for the prior belief, above which the CRA should be rewarded only if it announces the high rating and it is followed by success, and below which the CRA should be rewarded only if it announces the low rating. Notice that, quite intuitively, the CRA is never paid for announcing the high rating if it is followed by the project's failure. The idea behind the proof is that the CRA should, based on the prior, be paid for the event whose occurrence is the most consistent with its exerting effort, i.e., the one with the highest likelihood ratio.

The next proposition summarizes several properties of the Pareto frontier which will be important for our subsequent analysis.

Proposition 2 (Pareto Frontier) *Suppose the project is financed only after the high rating.*

- (i) *There exists v^* such that for all $v \geq v^*$ $e(v) = e^*$, but $u(v) < 0$.*
- (ii) *There exists $v_0 > 0$ such that (2) is slack for $v < v_0$ and binds for $v \geq v_0$. Moreover, $e(v_0) > 0$.*
- (iii) *Effort and total surplus are increasing in v , strictly increasing for $v \in (v_0, v^*)$.*

Part (i) of the proposition says that there exists a threshold promised value, v^* , above which the first-best effort is implemented. However, the resulting profit to the firm is strictly negative, violating individual rationality, and so this arrangement cannot be sustained in equilibrium. It will be handy to denote the highest promised value that can be delivered to the CRA without leaving the firm with negative profits by $\bar{v} \equiv \max\{v | u(v) = 0\} < v^*$.

There is an interesting economic reason why implementing the first-best effort requires the firm's profits to be negative. Suppose for concreteness $\gamma \geq \hat{\gamma}$ (the other case is similar), so that the CRA gets paid after history $h1$. Then the intuition is as follows. When effort is observable, the problem can be recast as saying that the firm chooses to acquire information itself rather than delegating this task to the CRA. But when the firm is making the effort choice it accounts for two potential effects of increasing effort. One benefit is the increased *probability* that a surplus is generated. The other is that investors will lower the interest rate to reflect a more accurate rating, leading to an increase in the *size* of the surplus. When the CRA is doing the investigation, and effort is not observed, it internalizes the fact that more effort generates a higher probability of the fee being paid. But the CRA cannot get a higher fee based on higher effort. So the only way to induce the CRA to exert

the first-best level effort is to set an extraordinarily generous fee that leaves the firm with negative profits.¹⁷

Part (ii) identifies the lowest value that can be delivered to the CRA on the Pareto frontier. This value, denoted by v_0 , is strictly positive. So the rating agency will still be making profits and will exert positive effort. It immediately follows from (ii) that for $v \leq v_0$ $u(v)$ does not depend on v and hence is constant; while if $v > v_0$, constraint (2) binds which means that $u(v)$ must be strictly decreasing in v .

Finally, part (iii) shows that the higher the CRA's profits, the higher the total surplus, and the higher the effort. This is an important result, and will be crucial for our further analysis. Intuitively it follows because unobservability of effort leads to its under-provision. To implement the highest possible effort, one needs to set fees as high as possible, extracting all surplus from the firm and giving it to the CRA. However, as part (i) tells us, implementing the first-best level of effort would result in negative profits to the firm. Combining (i) and (iii) tells us that the effort that can be implemented is strictly smaller than the first-best level.

Notice also that the firm's profits are maximized at v_0 . (This follows immediately from part (ii) of Proposition 2.) Thus the firm prefers a less informative rating than is socially optimal (as effort at v_0 is lower than that at \bar{v} or v^*), but the firm still prefers to have a informative rating (because effort is positive at v_0).

The function $u(v)$ is graphed in Figure 2. Recall that $u(v)$ only describes the part of the Pareto frontier which corresponds to the situation when the project is financed after the high rating and not financed after the low rating. The whole Pareto frontier is given by $\max\{-v, -1 + \pi_1 y - v, u(v)\}$, and the corresponding total surplus is $\max\{0, -1 + \pi_1 y, v + u(v)\}$.

To summarize, the fact that the CRA chooses its effort privately and has incentives to shirk has the following implications. First, the optimal compensation must involve outcome-contingent fees, which can be interpreted as rewards for establishing a good reputation. Second, the CRA exerts less effort, and hence there are more rating errors compared to the case when the CRA's effort is observable. These results are general — they do not depend on who orders a rating, and they will also hold in the extensions of the basic model that

¹⁷Formally, the firm's problem in the first best is $\max_e -\psi(e) + \pi_{h1}(e)(y - R(e))$, where the interest rate $R(e)$ solves the investors' break even condition $-\pi_h(e) + \pi_{h1}(e)R(e) = 0$. This implies $1/R(e) \equiv \pi_{1|h}(e) = \pi_{h1}(e)/\pi_h(e)$, the conditional probability of success given the high rating, will be strictly increasing in effort. The CRA's problem is $\max_e -\psi(e) + \pi_{h1}(e)f_{h1}$, where f_{h1} does not depend on e . Thus, in order to induce e^{FB} , f_{h1} must exceed $y - R(e)$, leaving the firm with negative profits: $\pi_{h1}(e)(y - R(e) - f_{h1}) < 0$.

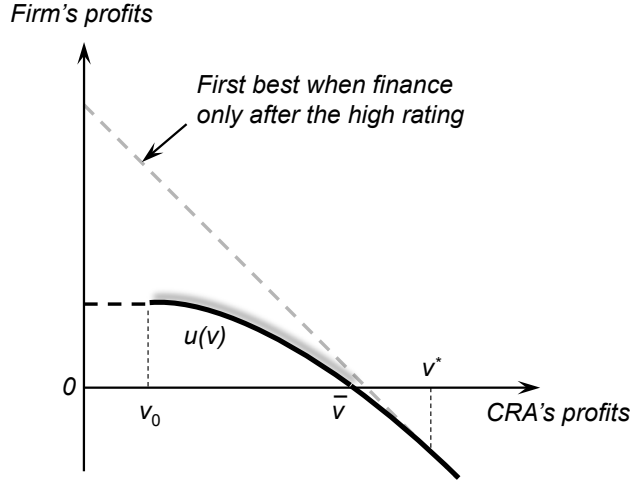


Figure 2: The Pareto frontier (the shaded area of the $u(v)$ curve) when the project is financed only after the high rating.

we will consider in Section 5.

Next, we consider how the equilibria will vary depending on who (X) orders the rating. For each X , the equilibrium will correspond to a different point $(v, u(v))$. Moreover, whether X even chooses to order a rating can differ across agents.

4.1 Social Planner Orders a Rating

Let us start with the case where the planner gets to decide whether a rating is ordered. Recall that we are considering equilibria where the total surplus is maximized. It immediately follows from Proposition 2 that the planner will choose the point $(\bar{v}, u(\bar{v}))$ on the frontier. This corresponds to maximum feasible CRA profits and effort, and zero profits for the firm. The implemented effort, which we denote by e^{SB} (where SB stands for the second best), is strictly smaller than e^{FB} . We summarize these results in the following proposition.

Proposition 3 ($X = \text{Planner}$) *If the social planner is the one who decides whether a rating should be ordered, then*

- (i) *The maximum total surplus in equilibrium is $S^{SB} = \max\{0, -1 + \pi_1 y, \bar{v} + u(\bar{v})\}$;*
- (ii) *$e^{SB} \leq e^{FB}$, $S^{SB} \leq S^{FB}$, with strict inequalities if $e^{FB} > 0$.*

Figure 3 uses a numerical example to compare the total surplus and effort in the first- and second-best (i.e., when the planner orders a rating) cases as functions of γ (depicted




Figure 3: The total surplus (left) and effort (right) as functions of the prior belief γ .

with solid black and solid gray lines, respectively). The thin dotted line in the left panel is $-1 + \pi_1 y$, the total surplus if the project is financed without a rating. The total surplus if the project is not financed without a rating is zero. Therefore, the total surplus in the first-best case, S^{FB} , is the upper envelope of three lines, 0, $-1 + \pi_1 y$, and $v^* + u(v^*)$. Similarly, the total surplus in the second-best case, S^{SB} , is the upper envelope of 0, $-1 + \pi_1 y$, and $\bar{v} + u(\bar{v})$.

From Figure 3 it is apparent that the planner decides not to acquire a rating for some values of γ when one would be acquired if effort were observable. The reduced propensity to get the rating occurs because the total surplus from acquiring the rating is lower. Thus, graphically, the interval on which the upper envelope of the three lines equals $\bar{v} + u(\bar{v})$ is smaller than that in the first-best case.

4.2 Issuer Orders a Rating

Next, we consider the case where the firm is the one who decides whether to order a rating. Recall that the CRA sets the fee schedule, and hence it will post the highest fee that the firm is willing to pay. The firm's willingness to pay equals its profit if it chooses not to order a rating. Without a rating, investors finance the firm's project if and only if $-1 + \pi_1 y > 0$. Since investors break even, the firm's profit in this case is $\underline{u} \equiv \max\{0, -1 + \pi_1 y\}$. Thus, if a rating is acquired in equilibrium, the firm receives \underline{u} , and the corresponding value to the CRA is $v^{iss} \equiv \max\{v | u(v) = \underline{u}\} \leq \bar{v}$, with strict inequality if $-1 + \pi_1 y > 0$ since $u(v)$ is strictly decreasing in v for $v > v_0$. Denote the total surplus and effort in the issuer-pays

case by S^{iss} and e^{iss} , respectively. Recall from Proposition 2 that the total surplus and effort are increasing in v . This leads us to the following result:

Proposition 4 ($X = \text{Issuer}$) *Suppose the firm decides whether to order a rating. Then*

- (i) *The maximum total surplus in equilibrium is $S^{iss} = \max\{0, -1 + \pi_1 y, v^{iss} + u(v^{iss})\}$;*
- (ii) *$e^{iss} \leq e^{SB}$, $S^{iss} \leq S^{SB}$, with strict inequalities if $e^{SB} > 0$ and $-1 + \pi_1 y > 0$.*

As usual, the firm will decide not to ask for a rating if the prior belief γ is sufficiently close to zero or one. Moreover, since the implemented effort with the firm picking whether to request a rating is lower relative to when the planner picks, rating acquisition will occur on a smaller set of priors in the former case than in the latter.

The total surplus and implemented effort in the case when the issuer orders a rating are depicted with dashed gray lines on Figure 3. As described in Proposition 4, when $-1 + \pi_1 y > 0$, the total surplus and effort are lower than when the planner orders a rating. Notice that e^{iss} decreases with γ when $-1 + \pi_1 y > 0$ because the firm's outside option is $-1 + \pi_1 y$, and π_1 increases with γ .

To summarize, the issuer-pays model leads lower rating precision and total surplus than the planner would attain, because the option of receiving financing without a rating reduces the firm's willingness to pay for a rating. That is, our model predicts that the issuer-pays model is indeed associated with more rating errors than is socially optimal. As we will see in the next section, the rating errors are also larger compared to the investor-pays model.

4.3 Investors Order a Rating

Consider finally the case when each investor decides whether to order a rating. We will show that this case results in a lower total surplus relative to the planner's case because investors are competing over financing terms that they offer to the firm conditional on the rating. As we will see, the comparison of the total surplus and effort relative to the issuer-pays case will depend on the prior γ .

The following proposition identifies the total surplus in this case.

Proposition 5 ($X = \text{Investors}$) *Suppose each investor decides whether to order a rating. Then $S^{inv} = \max\{0, v^{inv} + u(v^{inv})\}$, where $v^{inv} = \bar{v}(= v^{iss})$ if $-1 + \pi_1 y \leq 0$, and $v^{inv} \in (v^{iss}, \bar{v})$, otherwise.*

Notice the differences between the expressions for S^{inv} , S^{SB} in Proposition 3, and S^{iss} in Proposition 4. When investors pay, the term $-1 + \pi_1 y$ does not appear in the problem. For γ sufficiently close to one, it is socially optimal not to ask for a rating and always finance the project, so that S^{SB} (and S^{iss}) equal $-1 + \pi_1 y$. Investors, however, will choose to ask for a rating even when it is inefficient.¹⁸

The intuition is as follows. If the project is financed without a rating, then all surplus from the production ($-1 + \pi_1 y$) goes to the firm, while the CRA earns nothing. The CRA can try to sell a rating; it would not succeed if the planner controls whether it should be ordered, unless the generated surplus is at least $-1 + \pi_1 y$ (or unless the firm's profit is at least that amount, in case the firm orders a rating). However, when investors order a rating, they are not concerned with the total or the firm's surplus. They make zero profits, and they can always pass along the costs of getting a rating to the firm, while the CRA generates profits.

But why do investors necessarily choose to order a rating if they earn zero profits either way? To show that this must be the case, we prove that if no one asked for a rating, then one investor could generate profits by ordering a rating, hiding it from other investors, only investing if it is high, but charging the same or a slightly lower rate of return as other investors. Knowing this, the CRA can set fees low enough to entice someone to order a rating and hence break any equilibrium where no one is ordering a rating.

The other difference in the expressions for surpluses is that the promised value to the CRA in the investor-pays case is v^{inv} , which lies in between v^{iss} and \bar{v} , strictly whenever it is optimal to finance the project ex-ante, i.e., when $-1 + \pi_1 y > 0$. Therefore by Proposition 2, the implemented effort (and hence the rating precision) if investors ask for a rating is lower than if the planner asks for a rating, but higher than if the issuer does. The reason for $v^{inv} < \bar{v}$ is that the option to finance without a rating caps interest rates, and therefore caps fees that investors are willing to pay to the CRA. (This interest rate cap is $R = 1/\pi_1$, which solves $-1 + \pi_1 R = 0$.) And $v^{inv} > v^{iss}$ because the firm pays the same rate of return to investors as if there was no rating (R , defined above), but receives financing less often — only when the rating is high (without a rating, it would be financed with probability one). Hence the firm's profits are lower when the investor-pays than when the issuer does, $u(v^{inv}) < u(v^{iss})$, which in turn implies that $v^{inv} > v^{iss}$.

The total surplus and effort in the case when investors order a rating are plotted with dashed-dotted black lines on Figure 3. As one can see, the comparison between the total

¹⁸ Notice that in this case equilibrium payoffs actually lie inside the (constrained) Pareto frontier.

surplus in the issuer-pays and investor-pays cases depend on the prior belief about the project's quality. When the project is not profitable to finance ex-ante, i.e., when $-1 + \pi_1 y \leq 0$, the total surplus and effort in both models are equal, and coincide with what the planner achieves. However, when $-1 + \pi_1 y > 0$, the issuer-pays model leads to a lower total surplus than the investor-pays model for intermediate values of γ , but performs better if γ is sufficiently high. Also note that e^{inv} decreases with γ when $-1 + \pi_1 y > 0$, because $R = 1/\pi_1$ decreases with γ , which means the fees that investors pay to the CRA are falling as γ rises.

The following result formally summarizes the comparison of the total surplus and effort in the different models.

Corollary 1 (i) *If $-1 + \pi_1 y \leq 0$ then $S^{inv} = S^{iss} = S^{SB}$ and $e^{inv} = e^{iss} = e^{SB}$.*

(ii) *Suppose $-1 + \pi_1 y > 0$. Then*

(a) *$e^{inv} < e^{SB}$ if $e^{SB} > 0$, and $e^{inv} > e^{SB}$ if $e^{SB} = 0$;*

(b) *$e^{inv} \rightarrow 0$ and $S^{inv} \rightarrow S^{SB}/2$ as $\gamma \rightarrow 1$;*

(c) *$S^{inv} < S^{SB}$;*

(d) *$e^{iss} < e^{inv}$, with strict inequality if $e^{inv} > 0$;*

(e) *$S^{iss} < S^{inv}$ if $S^{inv} \geq -1 + \pi_1 y$, and $S^{iss} > S^{inv}$ otherwise.*

To summarize, the investor-pays model yields higher rating accuracy than the issuer-pays model, but lower than under the planner. The reason is that investors do not care about the firm's outside option, but the option to finance without a rating caps interest rates, and hence fees. On the other hand, investors ask for a rating too often, even when it is socially inefficient to do so.

Notice that the assumption that investors who do not pay for a rating can be excluded from learning it (directly, or through observing contract terms offered to the firm by informed investors) plays an important role. If this is not the case and the spread of information cannot be precluded, investors will want to free-ride on others paying for a rating. As a result, no rating will be acquired in equilibrium, and investors will make their financing decisions solely based on the prior. Until the mid 1970s, the investor-pays model was widely used. However, the rise of photocopying made protecting the sort of information described above became increasingly impractical, which arguably resulted in the switch to the issuer-pays model.

5 Extensions

We now consider four variants of the baseline model. Our first extension explores the effect of allowing more than one rating agency. Next, we consider the implications of allowing the CRA to misreport its information. Third, we look at differences in ratings for securities which differ in their ease of monitoring. The last modification introduces a second period in the model so that the propensity to downgrade a security can be studied.

5.1 Multiple CRAs

If multiple ratings are acquired in equilibrium, the problem becomes quite complicated. In particular, contracts will depend on CRAs' relative performance (i.e., a CRA's compensation would in part depend on other CRAs' ratings). Moreover, if ratings are acquired sequentially and are only published at the end, in the issuer-pays model the firm's decision whether to acquire the second rating will depend on its first rating. Since this rating is the firm's private information, it introduces an adverse selection problem. The analysis of this problem is sufficiently complicated that we leave it for future research.

Instead, as a first step, we restrict our attention to the case when, even though there are multiple rating agencies, only one rating is acquired in equilibrium. (Of course, this may or may not happen in equilibrium, so we simply operate under the assumption that it does.)

We modify the timing of our original model as follows. The game starts by CRAs simultaneously posting fees. The issuer then chooses which CRA to ask for a rating.¹⁹ Under these assumptions the problem becomes very simple to analyze. CRAs compete in fees, which leads to maximizing the issuer's profits. Recall from Proposition 2 that the firm's profits are maximized at v_0 .²⁰ Hence, the total surplus in this case, denoted by S_{many}^{iss} , equals $\max\{0, -1 + \pi_1 y, v_0 + u(v_0)\}$. Let e_{many}^{iss} denote the corresponding level of effort. Since $v_0 < v^{iss}$, it immediately follows from part (iii) of Proposition 2 that $e_{many}^{iss} \leq e^{iss}$ and $S_{many}^{iss} \leq S^{iss}$, with strict inequalities if $e^{iss} > 0$.

We find this extension interesting because it suggests that a battle for market share and

¹⁹We will only analyze the extension to multiple CRAs for the issuer-pays arrangement.

²⁰Of course, now there are more players in the game. If there are N CRAs and the firm randomizes between whom to ask for a rating if it is indifferent, then each CRA receives v_0/N in expectation. The frontier on Figure 2 shows the surplus division between the CRA whose rating is ordered and the firm after the outcome of the randomization is observed, with other CRAs (as well as investors) receiving zero profits.

desire to win business will lead to lower fees, which means less accurate ratings and lower total surplus. However, the firm’s surplus is higher despite the lower overall surplus. Also note that despite the competition, the CRAs still make positive profits, as $v_0 > 0$.

5.2 Misreporting a Rating

We next return to our original model with one CRA. So far we assumed that the CRA cannot misreport its signal; now we relax this assumption and suppose that the CRA can lie. In addition to moral hazard, this creates an adverse selection problem. Solving for the optimal contract requires imposing additional constraints to our optimal contracting problem (1)–(4).

It is easy to check that if the CRA intends to lie, the most profitable way to do so is a double deviation: exert no effort, and always report whatever rating yields the highest expected fee. This should not be surprising because if the CRA intends to misreport, exerting effort is wasteful. Hence, the additional constraint that needs to be imposed in order to deliver a truthful report is

$$-\psi(e) + \pi_{h1}f_{h1} + \pi_{\ell}f_{\ell} \geq \max\{f_{h1}, f_{\ell}\}, \quad (5)$$

where we have already incorporated the result of Proposition 1 that $f_{h0} = 0$.²¹ The left-hand side of the above constraint shows the CRA’s payoff if it exerts effort and truthfully reports the acquired signal. The right-hand side is the value from exerting no effort and always reporting the rating that delivers the highest fee (or randomizing between the two, if the fees are the same).

Lemma 2 *Constraint (5) binds so long as $e > 0$.*

Recall from Proposition 1 that when the CRA cannot misreport its signal, only one of the two fees, f_{h1} or f_{ℓ} , is strictly positive. It immediately follows from Lemma 2 that with the possibility of misreporting, both f_{h1} and f_{ℓ} must be strictly positive.

Lemma 2 also implies that the total surplus generated if the rating is ordered when the CRA can lie is lower than in the case when it cannot lie. Also, the range of priors for which the rating will be ordered (by any agent) is smaller than when the CRA cannot lie. This is not surprising since essentially the option to lie gives the CRA leverage that allows it to

²¹It is easy to check that imposing the additional incentive constraint does not change this prediction.

extract fees in order to tell the truth. These fees were previously unnecessary and mean that the agents now become more cautious about using the CRA.

5.3 New Securities

Suppose some types of investment projects are inherently more difficult for the CRA to evaluate — presumably because they have a short track record that makes comparisons difficult, and there is no adequate rating model that has been developed yet. One way to model this in our framework is to parametrize the cost of effort as $\psi(e) = A\varphi(e)$, with $A > 0$, and think of a new type of security as the one with a higher value of A .²² A higher value of A means that it is more costly for a CRA to obtain a rating of the same quality for a new security.

Suppose that A increases to A' . We consider two scenarios. The first scenario is that the shift in A is anticipated, and fees for ratings change appropriately. Claim 1 in the Appendix shows that it is optimal to implement lower effort with A' than with A , which results in larger rating inaccuracies. Intuitively, since the marginal cost of information acquisition is higher, it is optimal to reduce effort.²³

The second scenario is that a reduction in A is unanticipated. In this case, fees remain unchanged. Suppose that the CRA can misreport its rating. Claim 2 in the Appendix shows that in this case constraint (5) with A' instead of A becomes violated (recall from Lemma 2 that it was binding with A). Thus, if the rating agency realizes that A is low once it starts evaluating the security, its optimal response is to exert zero and always report either h or ℓ , depending on the prior.

Thus, our model predicts that under both scenarios the quality of ratings deteriorates for new securities.

5.4 Delays in Downgrading

Suppose finally that there are two periods. The project requires investment in both periods, and the project quality is the same in both periods. The CRA exerts effort in each period

²²We do assume that everything else, in particular, parameters p_b , p_g , y , and γ remain the same.

²³The result that information acquisition is decreasing in the cost parameter is also obtained in Opp, Opp, and Harris (2012). However, in their case this result is obvious since the CRA can commit to any level of effort, and will choose less effort if its marginal cost is higher. In our model, the result is less straightforward since fees are optimally chosen, but nonetheless the new optimal fee structure results in lower effort.

to rate the project. In the optimal contract, all payments to the CRA will be made at the end of the second period, conditional on the history. Denote these payments by $f_{i,j}$, where $i, j \in \{h1, h0, \ell\}$, whenever effort is exerted in both periods.

Suppose the CRA announced the high rating in period 1, which was followed by the project's failure. After such a history, there might be a need to downgrade the security (i.e., announce the low rating in period 2). By the same argument as in Proposition 1, to provide incentives for the second period effort at this point the CRA should be paid either $f_{h0,h1} > 0$ or $f_{h0,\ell} > 0$. However, it is easy to show that the best way to provide incentives for effort in period 1 is to pay $f_{h1,h1}$ or $f_{\ell,\ell}$ rather than $f_{h0,h1}$ or $f_{h0,\ell}$. That is, in order to create powerful incentives for the initial rating, the contract should reward the CRA for getting a correct rating that stays the same in both periods. But if a mistake is made, it becomes optimal to create incentives for the second period to recognize the mistake and possibly change the rating.

This intuition suggests that there is a trade-off between providing incentives for effort in period 1 (the initial rating) and effort in period 2 after a 'mistake' (when performance did not match the rating). The optimal contract is designed to balance this trade-off. The desire to support effort in period 1 makes fees in period 2 after a mistake too low. This in turn makes period 2 effort too low, ex post. As a result, the probability of not downgrading conditional on the project quality being bad is too high ex post. Hence, the CRA will appear too slow to acknowledge mistakes. Remarkably, this inertia seems to be a very general property of an optimal compensation scheme.

6 Conclusions

We develop a parsimonious optimal contracting model that addresses multiple issues regarding rating performance. We show that when the CRA's effort is unobservable, a rating is less precise, and is acquired less often (on a smaller set of priors) than in the first-best case. Giving all surplus to the CRA maximizes rating accuracy and total surplus.

Regarding the question of pros and cons of the issuer- and investor-pays model, we find that in the issuer-pays model the rating is less accurate than in the second-best case. The reason is that the option to finance without a rating puts a bound on the firm's willingness to pay for one. The investor-pays model generates a more precise rating than the issuer-pays model, although still not as high as what the planner could attain. However, investors tend to ask for a rating even when it is socially inefficient, in particular, when the prior

about the project quality is sufficiently high. In addition, the investor-pays model suffers from a potential free-riding problem, which can collapse security rating all together.

We show that battle for market share by competing CRAs leads to less accurate ratings, which yields higher profits to the firm. We also find that rating errors tend to be larger for new securities. Finally, we demonstrate that optimal provision of incentives for initial rating and revision naturally generates delays in downgrading.

While we view the mileage that is possible with our very parsimonious framework as impressive, there are many ways in which the modelling can be extended. Perhaps most natural would be to allow the firm to have superior information about its investment opportunities relative to other agents. While a general analysis of moral hazard combined with adverse selection is typically quite complicated, there are a few things we can see in some interesting special cases.

First, suppose as in our original model that the firm has no initial wealth. In this case the only (net) payment that the firm can possibly make occurs when project succeeds, and either both types will want to make such a payment, or neither will. Thus only a pooling equilibrium exists, and the analysis is essentially the same as in our original model. By continuity, the same will be true if the initial wealth is positive but sufficiently small.

Alternatively, if the firm has sufficient internal funds (but not enough to fund the project), then it might be possible to separate firms with different information about their projects. Whether separation is possible will depend on investors' prior beliefs about the project. The prior matters because competing investors are offering contract terms to firms with private information about their projects, and as in Rothschild and Stiglitz (1976), this can lead to a breakdown of any separating equilibrium and equilibrium non-existence when there are sufficiently many good types.²⁴

In some cases, the presence of a rating rather than investors' contracts can be used to separate different types of firms. Consider, for example, a special case where the firm knows its type perfectly, and suppose X is the planner. Then there is a separating equilibrium where the CRA charges a flat fee for publishing an uninformative rating, investors finance any project that gets rated regardless of the rating, and only the good type gets rated. The total surplus in this case is the highest possible (and hence the first best is implemented),

²⁴The breakdown can occur because a deviating investor can offer a profitable pooling contract which will attract both types of firms. In general, pooling is never an equilibrium because a deviating investor can offer a contract that attracts only good types. So when there are sufficiently many good types, there may simply be no equilibrium.

because only good projects receive financing and the cost of effort is zero.²⁵ Sustaining such an equilibrium is not always possible, for example, in the issuer-pays model. The reason again is that the firm's outside option caps the fee that the firm is willing to pay for a rating.²⁶

A more interesting and also a more complicated case is when the firm has some private information about the project quality, but does not know it perfectly. In this case, even in a separating equilibrium a rating will no longer be fully revealing about the project's quality. Notice that in a separating equilibrium agents' prior beliefs about the quality of different types of firms will differ, and the CRA's precision in evaluating different types of firms can differ too. This means that the same signal for different types will lead to different posterior beliefs about project quality, which can be interpreted as receiving different ratings. Put differently, even though only two signals are possible, there are circumstances where the CRA will report more than two different ratings.

We leave a more complete treatment of this problem for future work.

²⁵The planner makes the firm pay for a rating out of its initial wealth. Formally, the fee f that makes it worthwhile for only the good type to participate, while the bad type opts out, must satisfy $p_b(y - R) \leq f \leq p_g(y - R)$, where $R = 1/p_g$ is the interest rate charged by investors.

²⁶Formally, when $-1 + \pi_1 y > 0$ it must be the case that $-1 + p_g y - f \geq -1 + \pi_1 y$, or $f \leq (p_g - p_b)y(1 - \gamma)$. If γ is close enough to one, then this constraint and $p_b(y - 1/p_g) \leq f$ cannot hold simultaneously.

A Appendix: Omitted Proofs

Proof of Lemma 1. The total surplus in the first-best case is $S^{FB} = \max\{0, -1 + \pi_1 y, \max_e -\psi(e) - \pi_h(e) + \pi_{h1}(e)y\}$, where the third term can be rewritten as $\max_e -\psi(e) + (1/2 + e)(-1 + p_g y)\gamma + (1/2 - e)(-1 + p_b y)(1 - \gamma)$. At $\gamma = 0$, the first term exceeds the other two terms: $0 > -1 + \pi_1 y = -1 + p_b y$ and $0 > \max_e -\psi(e) + (1/2 - e)(-1 + p_b y)$. At $\gamma = 1$, the second term exceeds the other two terms: $-1 + \pi_1 y = -1 + p_g y > 0$ and $-1 + p_g y > \max_e -\psi(e) + (1/2 + e)(-1 + p_g y)$. Hence at $\gamma = 0$ ($\gamma = 1$) it is optimal not to acquire a rating and never (always) finance the project.

Define γ^* such that $-1 + \pi_1 y = (-1 + p_b y)\gamma^* + (-1 + p_b y)(1 - \gamma^*) = 0$. We claim that at $\gamma = \gamma^*$, the third term exceeds the other two terms, and hence it is optimal to acquire a rating and only finance the project after the high rating. To see this, consider the first-order condition of the maximization problem in the third term,

$$\psi'(e) = (-1 + p_g y)\gamma - (-1 + p_b y)(1 - \gamma). \quad (6)$$

The right-hand side of this equation is strictly positive at $\gamma = \gamma^*$. Hence (6) has a unique solution $e > 0$ at γ^* . Moreover, it is always possible to obtain zero surplus by choosing $e = 0$. Since the problem is strictly concave in effort, $-\psi(e) - \pi_h(e) + \pi_{h1}(e)y$ must be strictly positive at the optimal e .

Next, we show that $\max_e -\psi(e) - \pi_h(e) + \pi_{h1}(e)y$ is strictly increasing and convex in γ . It will then follow that it must single-cross 0 at $\underline{\gamma} \in (0, \gamma^*)$ and $-1 + \pi_1 y$ at $\bar{\gamma} \in (\gamma^*, 1)$, proving the interval structure stated in the lemma. Indeed, by the Envelope theorem, $\partial[-\psi(e) - \pi_h(e) + \pi_{h1}(e)y]/\partial\gamma = \partial[(1/2 + e)(-1 + p_g y)\gamma + (1/2 - e)(-1 + p_b y)(1 - \gamma)]/\partial\gamma = (1/2 + e)(-1 + p_g y) - (1/2 - e)(-1 + p_b y) > 0$. Differentiating again yields $\partial^2[-\psi(e) - \pi_h(e) + \pi_{h1}(e)y]/\partial\gamma^2 = (-1 + p_g y - 1 + p_b y)\partial e/\partial\gamma = (-1 + p_g y - 1 + p_b y)^2/\psi''(e) \geq 0$, where the last equality follows from differentiating (6) with respect to γ . \square

Proof of Proposition 1. Let λ and μ denote the Lagrange multipliers on constraints (2) and (3), respectively. The first-order condition of problem (1)–(4) with respect to f_i , $i \in \{h1, h0, \ell\}$ is

$$(-1 + \lambda)\pi_i(e) + \mu\pi'_i(e) \leq 0, \quad f_i \geq 0,$$

with complementary slackness. Dividing by $\pi_i(e)$, one can see that the first-order condition which will hold with equality (resulting in the strictly positive fee) is the one that

corresponds to the highest likelihood ratio, $\pi'_i(e)/\pi_i(e)$. Straightforward algebra shows that

$$\frac{\pi'_{h1}}{\pi_{h1}} \begin{matrix} \geq \\ < \end{matrix} \frac{\pi'_\ell}{\pi_\ell} \Leftrightarrow \gamma \begin{matrix} \geq \\ < \end{matrix} \frac{1}{1 + \sqrt{\frac{p_h}{p_\ell}}},$$

$$\frac{\pi'_{h1}}{\pi_{h1}} > \frac{\pi'_{h0}}{\pi_{h0}} \quad \text{for all } \gamma,$$

which completes the proof. \square

Proof of Proposition 2. (i) Consider the case when $\gamma \geq \hat{\gamma}$ (the other case is similar). Define $f_{h1}^* = \psi'(e^*)/\pi'_{h1}(e^*)$ — the fee that implements e^* — and let $v^* = -\psi(e^*) + \pi_{h1}(e^*)f_{h1}^*$. Thus by construction e^* can be implemented at $v = v^*$. For $v > v^*$, it can be implemented by paying a history-dependent fee f_{h1}^* plus an upfront fee equal to $v - v^*$.

We now show that $u(v^*) < 0$. Since $u(v) = u(v^*) - (v - v^*)$ for $v \geq v^*$, it will follow that $u(v) < 0$ for $v \geq v^*$. We again will restrict our attention to the case of $\gamma \geq \hat{\gamma}$. Using (6), for the first-best effort to be implemented it must be the case that $\pi'_{h1}f_{h1}^* = (-1 + p_g y)\gamma - (-1 + p_b y)(1 - \gamma)$. Substituting this into the firm's payoff, obtain

$$\begin{aligned} u(v^*) &= -\pi_h + \pi_{h1}y - \pi_{h1}f_{h1}^* = (1/2 + e)(-1 + p_g y)\gamma + (1/2 - e)(-1 + p_b y)(1 - \gamma) \\ &\quad - [(1/2 + e)p_g \gamma + (1/2 - e)p_b(1 - \gamma)] \frac{(-1 + p_g y)\gamma - (-1 + p_b y)(1 - \gamma)}{p_g \gamma - p_b(1 - \gamma)} \\ &= \gamma(1 - \gamma)[p_b - p_g]y < 0, \end{aligned}$$

where the last equality follows from straightforward algebra.

(ii) We again focus on the case when $\gamma \geq \hat{\gamma}$ (the other case is analogous). Consider maximizing the firm's payoff while omitting constraint (2). Substituting from (3), the firm's payoff can be written as $-\pi_h(e) + \pi_{h1}(e)y - \pi_{h1}(e)f_{h1} = (-1 + p_g y)(1/2 + e)\gamma + (-1 + p_b y)(1/2 - e)(1 - \gamma) - \pi_{h1}(e)(\psi'/\pi'_{h1})'(e)$. The first-order condition with respect to effort is $0 = [(-1 + p_g y)\gamma - (-1 + p_b y)(1 - \gamma)] - \psi'(e) - \psi''(e)\pi_{h1}(e)/(p_b \gamma - p_g(1 - \gamma))$. The term in the square brackets is strictly positive, while the last two terms are zero at $e = 0$ by our assumptions $\psi'(0) = \psi''(0) = 0$. Thus $e = 0$ cannot maximize the firm's profits. The CRA's payoff is $\Pi(e) = -\psi(e) + \pi_{h1}(e)\psi'(e)/\pi'_{h1}(e)$, where $\Pi(0) = 0$. Differentiating yields $\Pi'(e) = \psi''(e)\pi_{h1}(e)/(p_b \gamma - p_g(1 - \gamma)) \geq 0$, with strict inequality for $e > 0$. Therefore, for v below some threshold value, which we denote by v_0 , (2) does not bind. Moreover, the optimal level of effort for $v \leq v_0$ is strictly positive.

(iii) For $v \leq v_0$, constraint (2) does not bind, and hence the total surplus and effort are

constant. For $v \geq v^*$, $e = e^{FB}$ and the total surplus equals S^{FB} . Suppose that $v \in (v_0, v^*)$. Then (2) holds with equality. Consider the case of $\gamma \geq \hat{\gamma}$ (the other case is analogous). Substituting from (3), constraint (2) can be written as $-\psi(e) + \pi_{h1}(e)\psi'(e)/\pi'_{h1}(e) = v$. Differentiating the left-hand side with respect to e yields $\pi_{h1}(e)\psi''(e)/\pi'_{h1}(e)$, which is strictly positive as $\psi''(e) > 0$ for $e > 0$ and $\pi'_{h1} > 0$ for $\gamma \geq \hat{\gamma}$. Thus the optimal choice of e must be strictly increasing in v . Since the total surplus $-\psi(e) - \pi_h(e) + \pi_{h1}(e)y$ is strictly increasing in e for $e < e^{FB}$, it follows that the total surplus is also strictly increasing in v . \square

Proof of Proposition 4. Part (i) is shown in the main text. Part (ii) immediately follows from part (iii) of Proposition 2 and the fact that $v^{iss} \leq \bar{v}$, with strict inequality if $-1 + \pi_1 y > 0$. Given our assumption that the firm can credibly announce that it did not get rated, it immediately follows that if $-1 + \pi_1 y > 0$, then $\underline{u} = -1 + \pi_1 y$, and thus $v^{iss} < \bar{v}$, as described in the main text. It turns out that this is still true even if the firm cannot credibly reveal to investors that it did not order a rating.²⁷ To see this, suppose to the contrary that the firm's payoff, \underline{u} , is strictly lower than $-1 + \pi_1 y$ when $-1 + \pi_1 y > 0$. Then one investor could promise financing to the firm regardless of a rating at the rate of return R such that $-1 + \pi_1 R = \varepsilon$, where $\varepsilon > 0$ is arbitrarily close to zero. By accepting this offer, the firm would generate profits equal to $\pi_1(y - R) = -1 + \pi_1 y - \varepsilon$, which is strictly higher than \underline{u} for ε low enough. The deviating investor receives profits equal to $\varepsilon > 0$, a contradiction. \square

Proof of Proposition 5. First, suppose that $-1 + \pi_1 y > 0$. We want to show that financing the project without a rating cannot happen in equilibrium. Suppose to the contrary that it does. In such equilibrium, the CRA and investors earn zero profits, while the firm captures all the surplus, $-1 + \pi_1 y$. Investors do not acquire a rating, always finance the project, and receive the (gross) rate of return R that solves $-1 + \pi_1 R = 0$. Suppose the CRA were to offer a flat fee f plus history-dependent fees f_{h1} if $\gamma \geq \hat{\gamma}$ and f_ℓ otherwise. The level of effort that these fees induce solves $f_{h1} = \psi'(e)/\pi'_{h1}(e)$ in the first case and $f_\ell = \psi'(e)/\pi'_\ell(e)$ in the second case. We want to show that it cannot happen that all investors choose not to order a rating if the fees are sufficiently small. Indeed, consider a deviation by one investor who orders a rating, only invests if the rating is high, and asks

²⁷In this case, the investors' contracts cannot distinguish between events when the rating has not been ordered, and when it has been ordered, was low, but the firm chose not to reveal it. Furthermore, if investors financed the project without a rating, then the firm with a low rating would choose not to announce it. Therefore the argument provided in the main text for showing that $\underline{u} = -1 + \pi_1 y$ when $-1 + \pi_1 y > 0$ does not work.

the issuer for the same, or a slightly lower, rate of return R as everyone else.

Consider the case $\gamma \geq \hat{\gamma}$ (the other case is analogous). Then after paying the flat fee, the deviating investor can generate profits equal to $\Pi(e) = -\pi_h(e) + \pi_{h1}(e)R - \pi_{h1}(e)\psi'(e)/\pi'_{h1}(e) = -\pi_h(e) + \pi_{h1}(e)/\pi_1 - \pi_{h1}(e)\psi'(e)/\pi'_{h1}(e)$, where $\pi'_{h1}(e) = p_h\gamma - p_\ell(1-\gamma)$. Notice that $\Pi(0) = 0$. Moreover, $\Pi' = -\pi'_h + \pi'_{h1}/\pi_1 - [\psi' + \psi''\pi_{h1}/\pi'_{h1}]$. Since $\psi'(0) = \psi''(0) = 0$, the second term is zero at $e = 0$, while straightforward algebra shows that the first term is strictly positive. Thus $\Pi'(0) > 0$, as the marginal cost of implementing an arbitrarily small strictly positive level of effort is zero, while the marginal benefit is positive. Thus the deviating investor can generate strictly positive profits by requesting a rating, and will agree to any strictly positive flat fee f that results in expected cost strictly lower than these profits. The CRA earns $f - \psi(e) + \pi_{h1}(e)\psi'(e)/\pi'_{h1}(e)$, where the last two terms go to zero as e goes to zero. Thus the CRA can sell a rating to investors by setting fees low enough.

To see that $v^{inv} < \bar{v}$ when $-1 + \pi_1 y > 0$, notice that there is an additional restriction because investors cannot charge the rate of return higher than R : $-1 + \pi_1 R = 0$. Indeed, suppose they charge $R_h > R$. Then there is a profitable deviation by one investor, namely, do not order a rating and offer $R' \in (R, R_h)$. The firm prefers R' to R_h , and the investor makes positive profits. Suppose that the restriction $R_h \leq R$ does not bind. Then maximizing the total surplus means pushing the firm's payoff to zero, $\pi_{h1}(y - R_h) = 0$, which implies $R_h = y > R$. A contradiction.

To see that $v^{inv} > v^{iss}$ when $-1 + \pi_1 y > 0$, notice that the firm pays the same rate of return $R_h = R$ as when it is financed without a rating, but it receives funding less often: $u(v^{inv}) = \pi_{h1}(y - 1/\pi_1) < \pi_1(y - 1/\pi_1) = u(v^{iss})$.

Suppose now that $-1 + \pi_1 y \leq 0$. We want to show that in this case $S^{inv} = S^{SB}$. Suppose first that if the planner is the one who orders a rating then asking for the rating and financing only after the high rating results in a negative total surplus. In this case, it is optimal not to ask for a rating and never finance, so that $S^{SB} = 0$. If investors are the ones who order a rating, then by definition $S^{inv} \leq S^{inv}$. Ordering a rating cannot be an equilibrium, since it would result in a negative payoff to at least one player. Hence in this case investors do not order a rating and never finance, so that $S^{inv} = S^{SB}$. Now suppose that $S^{SB} = \bar{v} + u(\bar{v})$. All surplus in the second best is captured by the CRA, and the firm and investors earn zero. Clearly, this is also an equilibrium when investors order a rating, and the one that maximizes the total surplus. Thus in this case $S^{inv} = S^{SB}$. \square

Proof of Corollary 1. The only part remains to be proven is (ii)-(b). We first show

that $e^{inv} \rightarrow 0$ as $\gamma \rightarrow 1$. Recall that investors charge $R_h = R = 1/\pi_1$. Their expected profits when $\gamma \geq \hat{\gamma}$ equal $0 = -\pi_h + \pi_{h1}/\pi_1 - \pi_{h1}f_{h1}$. As $\gamma \rightarrow 1$, $-\pi_h + \pi_{h1}/\pi_1 \rightarrow -(1/2 + e) + (1/2 + e)p_g/p_g = 0$. Therefore $f_{h1} \rightarrow 0$ as $\gamma \rightarrow 1$, and so $e^{inv} \rightarrow 0$. Since the project is financed only after the high rating, and the probability of this event goes to $1/2$ as $e \rightarrow 0$, it follows that $S^{inv} \rightarrow (-1 + \pi_1 y)/2 = S^{SB}/2$. At $\gamma = 1$, there is discontinuity in the total surplus, and $S^{inv} = -1 + \pi_1 y = S^{SB}$, as investors will not purchase an absolutely uninformative rating. \square

Proof of Lemma 2. Suppose to the contrary that (5) does not bind. Then the results in Proposition 1 apply, and $f_{h1} > 0 = f_\ell = f_{h0}$ if $\gamma \geq \hat{\gamma}$ and $f_\ell > f_{h1} = f_{h0}$ otherwise. Consider the case of $\gamma \geq \hat{\gamma}$ (the other case is analogous). Constraint (5) can then be written as $-\psi(e) + \pi_{h1}f_{h1} \geq f_{h1}$, or $-\psi(e) - \pi_{h0}f_{h1} \geq 0$. But the left-hand side is negative, strictly if $e > 0$. A contradiction. \square

Claim 1 *Suppose that $\psi(e) = A\varphi(e)$. Then the optimal level of effort in problem (1)–(4) strictly decreases with A .*

Proof. We use strict monotone comparative statics results from (Edlin and Shannon, 1998) to show that e is strictly decreasing in A . Define $a = 1/A$. Consider the case $\gamma \geq \hat{\gamma}$ (the other case is analogous). Using Proposition 2 and substituting from (3), problem (1)–(4) can be written as $\max_e -\pi_h(e) + \pi_{h1}(e)y - \pi_{h1}(e)\varphi'(e)/(a\pi'_{h1}(e))$ subject to $-\varphi(e) + \pi_{h1}(e)\varphi'(e)/\pi'_{h1}(e) \geq va$. Denote the objective function by $f(e, a)$. Differentiating, $f_e = -\pi'_h(e) + \pi'_{h1}(e)y - [\varphi'(e) + \pi_{h1}(e)\varphi''(e)/\pi'_{h1}(e)]/a$. Since $\varphi'(e) > 0$ and $\varphi''(e) > 0$ for $e > 0$, and $\pi'_{h1} > 0$ for $\gamma \geq \hat{\gamma}$, it follows that $f_{ea} > 0$. Next, differentiating the left-hand side of the constraint with respect to e , obtain $\partial[-\varphi(e) + \pi_{h1}(e)\varphi'(e)/\pi'_{h1}(e)]/\partial e = \pi_{h1}(e)\varphi''(e)/\pi'_{h1}(e) > 0$ for $e > 0$. Thus, the constraint can be written as $g(e) \geq va$, where g is a strictly increasing function, or, equivalently, $e \in \Gamma(a)$, where Γ is nondecreasing in a in the strong set order. Therefore the optimal choice of effort is strictly increasing in a , or strictly decreasing in A . \square

Claim 2 *Suppose f_{h1} and f_ℓ are the optimal choices of fees in problem (1)–(5), where $\psi(e) = A\varphi(e)$. If the CRA chooses effort facing such fees and $A' < A$, then (5) is violated, and hence the optimal response of the CRA is to exert zero effort and always report either h or ℓ , depending on γ .*

Proof. The CRA’s profits if it chooses to exert effort are $\pi(A) \equiv \max_e -A\varphi(e) + \pi_{h1}(e)f_{h1} + \pi_\ell(e)f_\ell$. By the Envelope theorem, $\pi'(A) = -\varphi(e) < 0$. Therefore the left-hand side of (5) with A' is strictly lower than that with A . Since the right-hand side of (5) does not change, and the constraint was binding with A , it now becomes violated. \square

References

- BAR-ISAAC, H., AND J. SHAPIRO (2010): “Ratings Quality Over the Business Cycle,” mimeo.
- BECKER, B., AND T. MILBOURN (2011): “How Did Increased Competition Affect Credit Ratings,” *Journal of Financial Economics*, 101, 493–514.
- BLACKWELL, D., AND M. GIRSHICK (1954): *Theory of Games and Statistical Decisions*. Dover Publications Inc., New York.
- BOLTON, P., X. FREIXAS, AND J. SHAPIRO (2012): “The Credit Ratings Game,” *Journal of Finance*, 67, 85–112.
- CHADE, H., AND N. KOVRIJNYKH (2012): “Information Acquisition, Moral Hazard, and Rewarding for Bad News,” mimeo.
- COVITZ, D., AND P. HARRISON (2003): “Testing Conflicts of Interest at Bond Rating Agencies with Market Anticipation: Evidence that Reputation Incentives Dominate,” Finance and Economics Discussion Series 2003-68. Board of Governors of the Federal Reserve System (U.S.).
- EDLIN, A., AND C. SHANNON (1998): “Strict Monotonicity in Comparative Statics,” *Journal of Economic Theory*, 81, 201–219.
- FAURE-GRIMAUD, A., E. PEYRACHE, AND L. QUESADA (2009): “The Ownership of Ratings,” *RAND Journal of Economics*, 40, 234–257.
- FINANCIAL CRISIS INQUIRY COMMISSION (2011): *Financial Crisis Inquiry Report: Final Report of the National Commission on the Causes of the Financial and Economic Crisis in the United States*. Government Printing Office, Washington DC.

- FULGHIERI, P., G. STROBL, AND H. XIA (2011): “The Economics of Solicited and Unsolicited Credit Ratings,” mimeo.
- GOEL, A., AND A. THAKOR (2011): “Credit Ratings and Litigation Risk,” mimeo.
- GROMB, D., AND D. MARTIMORT (2007): “Collusion and the Organization of Delegated Expertise,” *Journal of Economic Theory*, 137, 271–299.
- INDERST, R., AND M. OTTAVIANI (2009): “Misseling through Agents,” *American Economic Review*, pp. 883–908.
- (2011): “Competition through Commissions and Kickbacks,” *American Economic Review*, forthcoming.
- JIANG, J. X., M. H. STANFORD, AND Y. XIE (2012): “Does It Matter Who Pays for Bond Ratings? Historical Evidence,” *Journal of Financial Economics*, 105, 607–621.
- MASON, J., AND J. ROSNER (2007): “Where Did the Risk Go? How Misapplied Bond Ratings Cause Mortgage Backed Securities and Collateralized Debt Obligation Market Disruptions,” mimeo.
- MATHIS, J., J. MCANDREWS, AND J.-C. ROCHET (2009): “Rating the Raters: Are Reputation Concerns Powerful Enough to Discipline Rating Agencies?,” *Journal of Monetary Economics*, 56, 657–674.
- MEDVEDEV, A., AND D. FENNELL (2011): “An Economic Analysis of Credit Rating Agency Business Models and Ratings Accuracy,” Financial Services Authority Occasional Paper 41.
- MORGAN, D. P. (2002): “Rating Banks: Risk and Uncertainty in an Opaque Industry,” *American Economic Review*, 92, 874–888.
- MULLAINATHAN, S., AND A. SHLEIFER (2005): “The Market for News,” *American Economic Review*, 95, 1031–1053.
- OPP, C., M. OPP, AND M. HARRIS (2012): “Rating Agencies in the Face of Regulation,” *Journal of Financial Economics*, forthcoming.
- OWUSU-ANSAH, Y. (2012): “What Went Wrong? Examining Moodys Rated CDO Data,” mimeo.

- PAGANO, M., AND P. VOLPIN (2010): “Credit Ratings Failures and Policy Options,” in *Macroeconomic Stability and Financial Regulation: Key Issues for the G20* edited by Mathias Dewatripont, Xavier Freixas, and Richard Portes, 129–148.
- ROTHSCHILD, M., AND J. STIGLITZ (1976): “Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information,” *Quarterly Journal of Economics*, 90, 629–649.
- SANGIORGI, F., AND C. SPATT (2012): “Opacity, Credit Rating Shopping and Bias,” mimeo.
- SKRETA, V., AND L. VELDKAMP (2009): “Ratings Shopping and Asset Complexity: A Theory of Ratings Inflation,” *Journal of Monetary Economics*, 56, 678–695.
- UNITED STATES SENATE PERMANENT SUBCOMMITTEE ON INVESTIGATIONS (2011): “Wall Street and the Financial Crisis: Anatomy of a Financial Collapse,” Majority and Minority Staff Report.
- WHITE, L. (2010): “Markets: The Credit Rating Agencies,” *Journal of Economic Perspectives*, 24, 211–226.