

Mahmoud Elamin (USA)

Believe only what you can check: credit rating agencies, structured finance, and bonds

Abstract

This paper identifies rating verifiability as a key difference that explains why credit rating agencies (CRAs) failed to mitigate information asymmetries in the structured finance market but succeeded in the bond market. Two infinitely repeated models are analyzed. In the first, the rating is unverifiable, and there is no equilibrium where the CRA reveals its information. In the second, the rating is verified with some probability, and full information revelation is guaranteed for any verification probability, when the CRA is patient enough. The interaction between verification probability and CRA patience is also analyzed. It is recommended that investors should verify ratings of structured finance products more often to ensure truthful revelation by CRA.

Keywords: credit rating agencies, conflicts of interest.

JEL Classification: G20, D80.

Introduction

Prior to the recent financial crisis, CRAs were thought to mitigate information asymmetries in financial markets. They supposedly acted as independent certifiers of information, revealing all their information about rated projects' probabilities of default. Firms with projects to finance used the CRAs as a signaling device. Good firms asked to be rated, providing investors with independent certification of the quality of their projects. Trust in the CRAs' ability has not been ruined before the crisis. Although there had been some problems in rating bond issues in the past, but these "mishaps" were considered isolated incidents and did not shatter the markets' faith in the CRAs' ability. But since the financial crisis revealed an across-the-board failure in rating a whole class of assets-structured finance assets – confidence in the CRAs' ability to rate structured finance products has been severely damaged¹. Two questions born in the crisis motivate this paper: Why would the CRAs do a decent job of rating bonds, but fail in the rating of structured products? And what feature peculiar to structured finance products hinders the CRAs' ability to mitigate information asymmetries in this market?

This paper identifies a key difference between structured finance and bonds, and builds two infinitely repeated models to test its effects on the CRA's role: Bond ratings are verifiable, but structured finance ratings are effectively unverifiable². Two facts support this distinction. First, structured finance products are newer, they are more complex, and lack a long time-series on performance. This undermines investors'

confidence in analysts' opinions about CRA ratings. Second, the nature of structured finance is such that the information underlying their rating is less publicly available than for bonds. Structured finance products are securities that share three characteristics: Debt assets are pooled, liabilities backed by the debt assets are written, and the credit risk of the collateral asset pool is delinked from the credit risk of the originator³. Investors buy liabilities that are backed by the pool the originator puts together. The structuring process means that what is known about originator X does not help in evaluating structured products assembled by X ⁴. Evaluating these securities means evaluating the pool of loans these securities are written on. The public information about a particular issue is the information included in the issue's prospectus. Surprisingly, prospectuses of structured products frequently included scanty information about the exact composition of the pool of loans. Quite often that information was very general and aggregated. But without enough specific information about the pool publicly available, it is hard to verify the CRA's rating. By contrast, the bond market is different because bonds sold by a firm are generally backed by all the firm's assets and projects. Public information about the firm – its business projects, the quality of its assets, its income statement, balance sheet, and so on – can be used to assess the bond rating. Analysts use that public information to assess the bond issuer's probability of default and verify the CRA's rating.

The models I present show that rating verifiability is key to the proper functioning of CRAs. Models with insufficient rating verifiability show that the CRA does not fully disclose its information about the project's probability of default and therefore does

© Mahmoud Elamin, 2013.

The views expressed herein are those of the authors and not necessarily those of the Federal Reserve Bank of Cleveland or the Federal Reserve System.

¹ For example, among many others, Benmelech and Dlugosz (2009) discussed the "Credit Rating Crisis".

² As will become clear in the last section of this paper, effectively unverifiable means that these products either have a low probability of being verified or low confidence in the verification method.

³ Kendall and Fishman (2000) and Coval et al. (2009) also discussed this delinking.

⁴ Titman and Tsyplakov (2010) finds some evidence of the importance of originators in the commercial mortgage-backed securities market, which is a small subset of the structured finance market (less than 8% in 2011). Nonetheless, this information was not included in a typical prospectus, which is what investors had access to.

not fulfill its role. Evidence supporting this conclusion about the structured market includes the SEC's (2008) report, which depicted a CRA analyst worried that a structured product rating did not capture half of the deal's risk but that "it could be structured by cows and we would rate it." Another analyst lamented that the CRAs continued to create an "even bigger monster, the CDO market," adding "let's hope we are all wealthy and retired by the time this house of cards falters." The report also mentions the suspicious fact that rating agencies made "out-of-model adjustments" to structured finance ratings but did not document the rationale for those adjustments. Asked about the CRA failure in the structured market, former Treasury Secretary Henry Paulson said in 2010: "I don't want the rating agencies to be held up as the font of all truth." He said investors should "do some thinking for themselves". Models with enough verification show that there is an equilibrium where the CRA fulfills its role as the information asymmetry mitigator. This conclusion is supported by the descriptive literature. For example, White (2010) stated that: "Corporations and governments whose "plain vanilla" debt was being rated were relatively transparent, so that an obviously incorrect rating would quickly be spotted by others and would thus potentially tarnish the rater's reputation". Before structured finance growth exploded, Cantor and Packer (1994) noted that: "Analysts regularly make recommendations to buy and sell that implicitly confirm or contradict the agencies' ratings they provide alternative perspectives to the judgements of the rating agencies. The discipline provided by reputational considerations appear to have been effective, with no major scandals in the ratings industry".

The rest of the paper is organized as follows. Section 1 presents the results; section 2 review the literature; section 3 presents the one-shot game with no CRA; section 4 presents the unverified-rating, one-shot game with the CRA; section 5 presents the unverified-rating, infinitely repeated game; section 6 presents the infinitely repeated game with verification; and the final-section concludes.

1. Results

This paper starts with a one-shot game between investors with resources to lend, and firms that have access to a productive project but are resource-constrained. Information is asymmetric since firms know their projects' probability of default, but cannot credibly communicate it to investors. Investors compete by making offers that specify the issuance size and the rate of return they require. This game has no separating equilibrium where offers separate bad projects from good ones. But if an independent entity, a CRA, can observe the project's probability of

default, can it credibly reveal this information in equilibrium?

I consider two classes of models. The first, with unverifiable ratings, is similar to the structured finance market. In the other, the rating is exogenously verified with some probability. The model with verification has two interpretations: The first, where the verification probability is lower than a cutoff determined by the CRA's discount factor, still resembles structured finance. The second, where the verification probability is higher than the cutoff, resembles the bond market.

In the unverifiable ratings models, a CRA, fully informed of the probability of default when accessed, is added to the one-shot game. Informed firms decide to access the CRA or not. If accessed, the CRA makes a public announcement about the project type. Investors, seeing the access/no access decision and the CRA's subsequent announcement if there was access, Bertrand-compete given their posterior beliefs on default. Firms choose which offer to accept. No perfect Bayesian equilibrium (PBE), separating the types, exists. The informed CRA has no incentive to reveal its information when accessed.

If the game is infinitely repeated, would the continuation payoffs give the CRA an incentive to truthfully separate the types? Here, the CRA is the only strategic (long-run) player, with a discount factor. Firms and investors are short-lived players, and nature picks a new project every period. An informationally efficient equilibrium—a PBE with investors fully informed of the project type they face in every contingency with positive probability of occurring—does not exist, regardless of how high the CRA's discount factor is. Some equilibria, that are truthful only for a finite time, are also presented. The CRA fails to separate the types in both the one-shot game and the infinitely repeated game, when the rating is unverifiable.

Next, I consider the same infinitely repeated game, but with a twist. In every period when the project is rated, with some preset probability $0 < p \leq 1$ the rating is exogenously verified, the true default probability is publicly revealed, and investors can see whether the CRA's rating was truthful. I get two results here. The first (Proposition 5), shows that for every p , there exists a cutoff discount factor. When exceeded, an informationally efficient equilibrium exists. Second (Proposition 6), I fix the discount factor and look for the verification probability that ensures efficiency. I show another cutoff discount factor. If the CRA's discount factor is below the cutoff, there exists no informationally efficient equilibrium, even when verification occurs with probability one. Once it ex-

ceeds the cutoff, for every discount factor there exists a corresponding cutoff verification probability. If the rating is verified above that cutoff, there exists an informationally efficient equilibrium. If it is verified below that cutoff, no such equilibrium exists. To sum things up, if the CRA is not patient enough, there is no efficient equilibrium, even with certain verification. If the CRA is patient enough, verification has to be likely enough to guarantee efficiency.

2. Literature review

Before the recent crisis, the misconception about the similarity of bonds and structured finance from the CRAs' point of view was common. For example, BIS (2005) concluded that the possibility of "conflicts of interest" is no more severe in structured finance products than for single-name credit products, and that reputation mitigates any potential for bad behavior in these markets. Coval et al. (2009) noted that BIS (2005) articulated a widely-held view that market forces would solve potential problems.

Other recent descriptive work discussed features of structured finance and its ratings that support the various modeling assumptions of this paper. White (2010) compared structured finance to bonds: "...these mortgage-related securities were far more complex and opaque than were the traditional 'plain vanilla' corporate and government bonds, so rating errors were less likely to be quickly spotted by critics (or arbitrageurs)." Ashcraft and Schuermann (2008) implicitly mentioned nonverifiability of subprime MBSs: "Especially for investment grade ratings, it is very difficult to tell the difference between a 'bad' credit rating and bad luck"¹. SEC (2008) concluded that certain significant aspects of RMBS and CDOs ratings process were not always disclosed². Coval et al. (2009) showed by simulations the fragility of structured finance ratings relative to modest imprecision in the evaluation of underlying risks and exposure to systematic risks. Hence, a "slight" difference in two analysts' opinions concerning pool correlations might lead to widely differing opinions of a rating³.

Some empirical papers support the results of this paper. Covitz and Harrison (2003) analyzed bond rating migrations and concluded that rating changes do not appear to be importantly influenced by

CRAs' conflicts of interest. Griffin and Tang (2010) analyzed CDO ratings and showed that CRAs made positive CDO rating adjustments beyond what the CRAs' own models implied. Benmelech and Dlugosz (2009) examined Moody's structured and bond ratings and concluded that structured ratings were inflated. The magnitude of bond rating downgrades was relatively low and stable over time. In comparison, structured securities downgrades in the 2007-08 crisis were much more severe.

Some theory papers considered nonstrategic CRAs, ignoring CRA conflicts of interest. Skreta and Veldkamp (2009) considered a CRA that mechanically reveals the signal it gets about the true default probability and examined "rate shopping" in this context. If the issuer asked to be rated, then he chooses whether to reveal the signal or not. But if investors asked for the signal, it is always revealed. Carlson and Hale (2006) considered a global game with investors receiving private signals about creditworthiness. They treat the CRA as a public, noisy signal about creditworthiness. The authors assume full information revelation by the CRA and are only concerned with the effect of the CRA's public announcement on sovereign debt rollover.

Few theory papers have considered strategic CRAs, but in a one-shot game context. Lizzeri (1999) discussed information revelation and surplus extraction of a certification intermediary who commits to a fixed fee and disclosure code at the beginning of the game. Sellers decide to access the intermediary's service or not, and buyers move last. He found that a monopoly intermediary only reveals if the good has negative or positive value and captures a large share of the surplus. When intermediaries Bertrand-compete, their surplus drops to zero and full disclosure is provided. Kuhner (2001) presented a one-shot game with CRA payoffs determined by an exogenous "reputation function," supposedly standing for continuation payoffs in an unmodeled infinitely repeated game. He found an equilibrium with limited information revelation. Bolton et al. (2012) analyzed the interaction of rate shopping and competition in a one-shot game. Issuers are always rated but can choose to reveal the rating or not. Investors are either trusting and take the CRA's word at face value, or sophisticated and rationally update. The CRA has an exogenously specified reputation function (which is also analyzed in a two-period model) and loses its revenue if it lies. The authors assumed that the subsequent game between investors and issuers results in investors getting their reservation utilities. Importantly, they assume that the CRA's rating is verifiable. In the monopoly CRA case, they showed one equilibrium that exhibits either rating inflation or not, depending on the relative importance of the fee

¹ This supports the idea of low or no verifiability of structured finance products.

² Hence, hampering structured finance verifiability.

³ This is further supported by a quote from Mark Adelson's (director of structured finance research at Nomura Securities) testimony before the Committee on Financial Services on September 27, 2007, cited in Skreta and Veldkamp (2009): "The complexity of a typical securitization is far above that of traditional bonds. Despite the outward simplicity of credit-ratings, the inherent complexity of credit risk in many securitizations means that reasonable professionals starting with the same facts can reasonably reach different conclusions."

with respect to the reputation loss. They further considered a duopoly CRA market and did a welfare analysis on the effect of competition by comparing the equilibria they identified. Similarly, Sangiorgi and Spatt (2012) consider a one-shot game where an issuer can be pre-rated by CRAs (the indicative rating stage) and then decide to buy the rating or not. Unlike Bolton et al. (2012), investor here are always rational. They find that in the presence of transparency at the indicative rating stage, all ratings are disclosed and there is no bias. While in its absence, not all ratings are disclosed and rating bias arises.

Many authors and commentators believe that rate shopping exacerbated rating inflation in the structured finance market¹. But rate shopping is also prevalent in the bond rating market². Why is rate shopping potentially a problem in the structured finance market but not in the bond market? This paper provides a reasonable explanation: rating verifiability is what mitigates CRAs' conflict of interest in bond market.

The reputation loss function discussed in the literature can be endogenized in an infinitely repeated game with or without reputation. This paper analyzes an infinitely repeated game with no reputation. Mathis et al. (2009) and Elamin (2010) model reputation as incomplete information about the type of the long-run player, where one type always behaves in a preset way. The updated probability about his type is his reputation³. Rablen (2013), discussed later, extends Mathis et al. (2009). I explain below why reputation in itself is not important for resolving CRA conflict of interest.

Mathis et al. (2009) and Elamin (2010) independently considered a strategic CRA in an infinitely repeated game with reputation to explore whether reputation disciplines the CRA. Elamin (2010) added reputation to the model with the unverifiable ratings presented here.

With some probability, the CRA might be a truthful robot. I found that if there are only two projects, then when reputation is high enough, an informationally efficient equilibrium that separates project types exists. With three or more types, no matter how high reputation or the discount factor is, such an equilibrium does not exist. I concluded that reputation is not potent enough to ensure efficiency in the structured finance market.

Mathis et al.'s (2009) model is the closest to this paper, but with important differences. In Mathis et al.

(2009), firms are nonstrategic and are forced to access. Projects are deterministic: Bad projects always default, and good projects always pay⁴. The financing decision is discrete, in that investors either finance the project or not. Therefore, the "size of the investment" is fixed. Moreover, the CRA has two revenue sources. The first (structured finance revenue), depends on the CRA's reputation and actions. The second comes from "other sources," representing bond rating revenue among other things. The authors assumed this revenue is lost if the "structured market" reputation drops to zero. This assumption is crucial for their result that if the ratio of the revenue from other sources relative to structured finance is high enough, then a truthful equilibrium exists and reputation disciplines the CRA.

Rablen (2013) extends the Mathis et al. (2009) model by making the structured finance and bonds markets interdependent in a very particular way: the CRA has a "common reputation" affecting both markets, resulting in what he terms "reputational spill – over effects". Independently of this paper, he also identifies the difference in verifiability between the two markets, but only considers the effect of verifiability on the CRA's common reputation. He finds that when reputation is high, a patient enough CRA inflates structured finance ratings, but not bond ratings. Similarly to Mathis et al. (2009), the gist of the argument is that reputation disciplines the CRAs in markets with verifiability (bonds), but fail in markets without it (structured finance). Both Mathis et al. (2009) and Rablen (2013) argue that reputational concerns are what disciplines the CRA. Mathis et al. (2009) showed that in a reputation model with verifiability (when the project default shows the CRA lied), CRA truthfulness is guaranteed when it is patient enough. Rablen (2013) on the other hand, makes the verifiability more explicit and gets the same conclusion.

The models I present shed light on two problems with this result. First and more importantly, my paper shows that the crucial driving force behind the truthful equilibrium existence is not reputation, but verifiability. In the infinitely repeated game with no reputation presented here, verifiability still delivers CRA discipline. The project types assumed in Mathis et al.'s (2009) base model clearly identified if the CRA lied or not, and this verifiability is what delivered their truthful equilibrium. This is further confirmed by looking at Proposition 4 of Mathis et al. (2009), which shows that once project types did not allow verifiability, the CRA always lies. Rablen (2013) is more explicit about verifiability but still focuses on reputation as a disciplining force. This fo-

¹ Skreta and Veldkamp (2009) and Bolton et al. (2012) among others.

² For example, see Smith and Walter (2001, p. 12 and p. 22) and Cantor and Packer (1994).


³ Mailath and Samuelson (2006).

⁴ This is in the base model; they did consider extensions, and I discuss later how this supports my results.

cus on reputation is misplaced. Second, this paper argues that the bond market and the structured finance market are segregated because of the verifiability difference between the instruments involved. This goes against the logic of considering one CRA reputation for both markets. Moreover, it points out that a “reputational spill-over” effect is highly unlikely. This “spill-over” is assumed in both Mathis et al. (2009) and Rablen (2013).

One testable implication that distinguishes my theory from this reputation literature is whether bond rating

revenues have suffered after the crisis. CRAs’ reputation suffered a big blow during the crisis, but since truthfulness does not depend on reputations in bonds market, my theory predicts no loss in bond revenue. Figure 1 divides Moody’s revenue into structured finance revenue and sums up corporate, financial institution, and public ratings revenue into other rating revenue. It shows that although there was a substantial drop in structured finance rating revenue after the crisis, there was only a slight dip in other rating revenue during the crisis, with a sharp reversal and upward trend starting in 2008.



Note: Moody’s Annual Reports.

Fig. 1. Moody’s rating revenue (in millions of dollars)

3. One-shot game with no CRA

It is useful to examine the equilibria of the game without a CRA for three reasons. First, if these equilibria are “unintuitive”, then we should question the suitability of the chosen model. Second, this sheds some light on the viability of the proposed CRA role. The offers themselves could distinguish between the different borrowers. For example, a good borrower would pick one offer, the bad borrower would pick the other, and so on. This negates the need for a CRA. Third, the simple relationship between the investors’ offers and their beliefs about the project type shown here, is mimicked in the more complicated games of later sections and in the game of reputation of Elamin (2010).

3.1. Setup. This is a one-shot game between a borrower and two lenders who Bertrand-compete. A firm, with either a good or a bad project, is resource-constrained and must borrow to fund it. Both projects yield \bar{R} when they pay and zero in default and differ only in their probability of default. Projects are ex-post indistinguishable: it is not possible to perfectly distinguish between the two types ex post¹. Nature picks a project with default probability from the set $\{p_H, p_L\}$, with prior $0 < \eta < 1$ on p_L (a low quality project) and $0 < p_H < p_L < 1$. It reveals its choice only to the firm. The lenders are identical with log utility

and move simultaneously. Each investor i has an endowment of one unit of a good and chooses an offer (R_i, b_i) . He picks both how much to invest $0 \leq b_i \leq 1$, and the rate of return R_i in \mathcal{R} , and consumes $1 - b_i$ for sure. The firm sees both offers (R_1, b_1) and (R_2, b_2) ; and chooses which offer to accept, only when offers give different profits. Otherwise, nature randomizes picking each with equal probability.

When (R_i, b_i) is the selected offer, the firm earns zero profits if the project defaults (limited liability) and $(\bar{R} - R_i) b_i$ if it pays. An investor’s payoff depends on whether he is picked or not. The unpicked investor consumes his endowment and gets zero utility. The picked investor i consumes $1 - b_i$ when the project defaults, and $1 + (R_i - 1) b_i$ when it pays. Let $q = \eta p_L + (1 - \eta) P_H$ be the posterior on default when the prior is retained. And assume the following condition.

Condition 1. $q < \frac{\bar{R} - 1}{R}$.

Rewriting Condition 1 as $\bar{R}(1 - q) > 1$, shows that the project’s expected payment under the prior exceeds saving the endowment to the end of the period. It guarantees that investors are willing to invest some of their endowment under the prior.

To define strategies, let S be the space of offers that give the firm the same profits. $S = \{(R_1, b_1), (R_2, b_2)\} \in (\mathcal{R} \times [0, 1])^2 : (\bar{R} - R_1) b_1 = (\bar{R} - R_2) b_2\}$.

¹ This differs from Bolton et al. (2012) and the base model of Mathis et al. (2009) where a good project does not default.

On S , the firm does not move, and nature randomizes equally on both offers. The firm's strategy is $\sigma_{Firm} : \{p_H, p_L\} \times (\mathcal{R} \times [0, 1])^2 \times S \rightarrow \{1, 2\}$. When called upon to move, the firm chooses which offer to accept. Investor i 's strategy is $\sigma_i : \emptyset \rightarrow \mathcal{R} \times [0, 1]$. The perfect Bayesian equilibrium (PBE) definition follows.

Definition 1. A perfect Bayesian equilibrium (PBE) (σ, μ) is a collection of strategies $\sigma = (\sigma_1, \sigma_2, \sigma_{Firm})$ and beliefs $\mu = \eta$ s.t. when called upon to move, the firm picks the offer that yields the highest profits. Investor i , retaining his prior, acts optimally given σ_j and σ_{Firm} .

3.2. Equilibrium characterization. Let G be the set of solution offers (R, b) that solve problem P1: $\text{Max}_{R, b} (\bar{R} - R_1) b$ s.t. $0 < b < 1$ and $q \log(1 - b) + (1 - q) \log(1 + (R - 1) b) = 0$.

Proposition 1. Under Condition 1, if (σ, μ) with corresponding offers (R_1, b_1) and (R_2, b_2) is a PBE, then $\forall i (R_i, b_i) \in G$ and the firm picks the highest expected-profits offer if called upon to move.

Proof. Relegated to Appendix.

The intuition for Proposition 1 is clear. Investors Bertrand-compete, maximizing the firm's profits until they get the value of their outside option, which is consuming their endowment and netting zero. Proposition 1 showed that an offer (R_i, b_i) by investor i should solve problem P1. Lemma 1 shows the existence of a unique solution to P1. Hence, investors propose exactly the same offer when they move.

Lemma 1. Under Condition 1, there exists a solution (R, b) that solves P1. Moreover, the solution is unique and is characterized by the following two equations:

$$b = 1 - \left[\frac{q}{(1-q)(\bar{R}-1)} \right]^{1-q} \quad \text{and} \quad R = \frac{1}{q} \left[(1-q)\bar{R} - 1 \right] \frac{1-b}{b}.$$

Proof. Relegated to Appendix.

Under Condition 1, the unique equilibrium has the following properties. First, the optimal issuance size b decreases in the prior q : investors demand less issuance size when they perceive the project to be riskier. Second, optimal firm profits decrease in q : firms perceived to have riskier projects get lower profits. Hence, the good firm has an incentive to send a (costly) signal that it has a better project, later when a CRA with costly access is introduced. This section highlights the need for a CRA; without one, investors cannot distinguish between the types of borrowers. No separating equilibrium exists because when the firms rank offers according to the profits they yield, both types of the firm rank offers

in exactly the same order. This kills the capacity to give an offer tailored to a particular type.

4. One-shot game with a CRA

In this section, I introduce a CRA which is informed of the true project type when accessed. The good firm has an incentive to reveal its type since that yields it higher profits. Would the informed CRA fulfill its role in the market, revealing its information when accessed and mitigating the information asymmetry?

4.1. Setup. I only discuss the differences from the game of section 1 here. A CRA is added to the game, and the informed firm decides whether to access it or not. If accessed, the CRA is fully and costlessly informed of the true project type¹ and makes a public announcement from the set $\{H, L\}$. In return for its services, the CRA deducts a fee proportional to the selected issuance size and invests what remains. This payment scheme parallels reality in two ways: First, the payment is upfront² and does not depend on how the project fares. Second, the CRA's payment increases with the issuance size. Whether there is access and a subsequent CRA public announcement or not, investors Bertrand-compete, making offers according to their (common) posterior beliefs about the project's default probability. As before, firms choose which offer to accept, only when offers give different profits. When they do not, nature randomizes, picking each with equal probability.

The firm earns zero profit when the project defaults (limited liability). If the project pays and (R, b) is the selected offer, it earns $(\bar{R} - R) b$ if the CRA was not accessed, and $(1 - \epsilon) (\bar{R} - R) b$ if the CRA was accessed. The CRA gets zero when not accessed. If accessed, it gets payment ϵb , where $\epsilon > 0$ and small (see Condition 3), and b is the issuance size of the selected investor's offer after access. The investor's payoff depends on whether he is picked or not. The unpicked investor consumes his endowment and gets zero. If the CRA was not accessed, the picked investor with offer (R, b) consumes $1 - b$ when the project defaults, and $1 + (R - 1) b$ when the project pays. If the CRA was accessed, the picked investor consumes $1 - b$ when the project defaults, and $1 + (R(1 - \epsilon) - 1) b$ when the project pays. The timeline is illustrated in Figure 2.

¹ This setup accommodates an alternative interpretation where the CRA sees an imperfect signal instead of the true project type. p_L and p_H represent the CRA's posteriors after a signal is received.

² Upfront fee means it is deducted from the issuance size before the investment is made, and does not depend on how the project fares.




Fig. 2. Timeline of events in the one-shot game

The author now present two conditions under which the game becomes interesting.

Condition 2. $p_L < \frac{(1-\epsilon)\bar{R}-1}{(1-\epsilon)\bar{R}}$.

Condition 3. $(1-\epsilon)\bar{R} - \frac{((1-\epsilon)\bar{R}-1)^q}{(1-q)^{1-q}q^q} > \bar{R} - \frac{(\bar{R}-1)^{p_L}}{(1-p_L)^{1-p_L}p_L^{p_L}}$.

Condition 2 states that even when the project is thought to be bad and after adjusting for the fee, the project’s expected payment is better than just saving the endowment. Condition 3 states that the profit when the firm accesses the CRA and investors retain their priors is strictly higher than when the firm plays no access and investors believe it to be bad. Hence, the fee is small enough that firms prefer to access and pool, than not access and be considered bad. In particular, the condition implies that $\epsilon < 1$. Given Condition 3, we have that Condition 2 implies Condition 1. This fact will be useful when we analyze the contingency after no access.

The firm moves twice. First, knowing the project type, it randomizes on access and no access: Its first move: $\sigma^1_{Firm}: \{p_H, p_L\} \rightarrow \Delta\{A, NA\}$ and at the end of the game it moves again if the offers give it different profits. Let S be the set of offers that give same profits. The firm’s second move: $\sigma^2_{Firm} \{p_H, p_L\} \times \{(A, H), (A, L), NA\} \times \{(\mathcal{H} \times [0, 1])^2 \setminus S\} \rightarrow \{1,2\}$. Seeing the true project type, the CRA randomizes on announcements after p_H , and after p_L . $\sigma_{CRA}: \{p_H, p_L\} \rightarrow \Delta\{H, L\}$. Each investor i chooses an offer, given what he observes, $\sigma_i: \{(A, H), (A, L), NA\} \rightarrow \mathcal{H} \times [0, 1]$.

4.2. Perfect Bayesian equilibrium. A perfect Bayesian equilibrium is a collection of strategies and beliefs. Strategies are optimal given beliefs. On-path posteriors are derived from strategies using Bayes’s rule. Off-path, where Bayes’s rule does not apply, a PBE puts no restrictions on posteriors.

Definition 2. A perfect Bayesian equilibrium (PBE) is a pair (σ, μ) where $\sigma = (\sigma^1_{Firm}, \sigma_{CRA}, \sigma_1, \sigma_2, \sigma^2_{Firm})$ is a collection of strategies and $\mu = (\mu(A, H), \mu(A, L), \mu(NA))$

1 denotes the investors’ beliefs in each of their information sets about the probability of facing αp_L project in those contingencies. Satisfying two requirements:

- ◆ Given μ : σ^2_{Firm} is optimal σ_i is optimal given $(\sigma_j, \sigma^2_{Firm})$; σ_{CRA} is optimal at each node (after p_H and after p_L) given $(\sigma_1, \sigma_2, \sigma^2_{Firm})$; σ^1_{Firm} is optimal at each node given $(\sigma_{CRA}, \sigma_1, \sigma_2, \sigma^2_{Firm})$.
- ◆ μ is derived from Bayes’s rule when possible, with no restrictions where it does not apply.

Understanding how investors’ optimal offers depend on their beliefs simplifies the analysis. The contingency after no access is similar to the case without a CRA considered before, with beliefs determined by $\mu(NA)$. Investors’ optimal offers are determined by Proposition 1 and Lemma 1. In the contingencies after access, investor’s optimal offers are similar to before except for a slight alteration in the constraints corresponding to the fact that ϵ is deducted from the issuance size of the selected offer as the CRA’s payment. Two lemmas, detailing investors’ optimal offers given beliefs in the three possible contingencies, follow. The proofs follow closely the proofs of Proposition 1 and Lemma 1, and are omitted.

Lemma 2. Let (σ, μ) be a PBE with corresponding sets of offers (R_i, b_i) . Then under Condition 2 and Condition 3.

- ◆ In the contingency following access, $\forall i (R_i(A, .), b_i(A, .))$ solves Problem P2: $(1-\epsilon) \text{Max}_{R, b} (\bar{R} - R)b$ s.t. $0 < b < 1$ and $q \text{Log}(1-b) + (1-q) \text{Log}(1 + (R(1-\epsilon) - 1)b) = 0$, where $q(A, .) = \mu_i(A, .)p_L + (1 - \mu(A, .))p_H$. And the firm picks the offer with the highest expected profit if called upon to move.
- ◆ In the contingency following no access, $\forall i (R_i$

¹ As is standard (Fudenberg and Tirole, 1991), the PBE definition has only one system of beliefs instead of a system for each investor. On-path beliefs are identical since they satisfy Bayes’s rule, but off-path beliefs could possibly differ. If off-path beliefs differ however, there exists no strategies that along with the differing beliefs, form part of a PBE. Nothing is lost by sticking with one system of beliefs. This nonexistence issue is very similar in spirit to a nonexistence issue in a first-price auction, when ties are resolved by randomization.

- ◆ $(N A), b_i (NA))$ solves Problem P3: $\text{Max}_{R,b} (\bar{R} - R)b$ s.t. $0 < b < 1$ and $q \text{Log}(1 - b) + (1 - q) \text{Log}(1 + (R - 1)b) = 0$, where $q(NA) = \mu(NA) p_L + (1 - \mu(NA)) p_H$. The firm picks the offer with the highest expected profit if called upon to move.

I note here that Problem P3 is analogous to Problem P1; hence Lemma 1 ensures that a unique solution exists and provides the solution in terms of parameters. Note that Condition 2 and Condition 3 imply Condition 1 which appears in Lemma 1. I turn to problem P2 to show an analogue to Lemma 1, ensuring that a solution to problem P2 exists and is unique.

Lemma 3. *Under Condition 2 and Condition 3, there exists a solution (R, b) that solves P2. Moreover, the solution is unique and is characterized by the following two equations:*

$$b = 1 - \left[\frac{q}{(1-q)(1-\epsilon)R-1} \right]^{1-q} \text{ and } R = \frac{1}{q} \left[(1-q)\bar{R} - \frac{1}{1-\epsilon} \right] \frac{1-b}{b}.$$

In every contingency, both investors make the same offer, determined by the equations of Lemmas 1 and 3. The unique offer depends on the past only through beliefs. The crucial effect of the firm’s access and no-access decision, and the subsequent announcement, is determining investors’ beliefs. Once the beliefs are determined, investors’ optimal offers are uniquely determined. Moreover, since investors make the same offer in every contingency, the firm will not move a second time, and nature resolves the tie. To simplify the analysis, this second move is ignored from now on.

4.3. Informationally efficient equilibrium. An informed CRA was introduced, in the hope it will mitigate information asymmetries between borrowers and lenders. The notion of efficiency considered in this paper reflects that role, and an efficient equilibrium is called informationally efficient. Intuitively speaking, in an informationally efficient equilibrium, investors are always informed of the type of project, and there is no information asymmetry. A more formal definition follows.

Definition 3. *A PBE is informationally efficient if investors are correct in their beliefs about the project type on-path¹.*

In what follows, I will analyze what kind of role the CRA plays in any PBE.

¹ To prevent possible confusion, “off-path” refers to contingencies where Bayes’s rule does not apply. The definition of informational efficiency is less stringent than to require that investors are always informed of the true project type. It allows the possibility that investors’ off-path beliefs might not be correct. Since investors’ off-path information sets “do not occur”, and since PBE puts no restrictions on off-path beliefs, this more lenient requirement is more natural.

Proposition 2. *The one-shot game has no informationally efficient PBE. Moreover, whenever $\mu(A, H)$ and $\mu(A, L)$ are determined by Bayes’s rule, $\mu(A, H) = \mu(A, L)$.*

Proof. Assume there is an informationally efficient equilibrium. There are three possible cases to consider. In case 1, both firms access and the CRA reveals the true types. In case 2, the bad firm accesses and the good firm does not. That the bad firm never accesses if its true type is revealed after access, eliminates cases 1 and 2 as possible equilibria. The worst that could happen after no access is to be considered bad, but even then it saves the fee². For case 3, assume the good firm accesses and the bad firm does not. This will work only when the CRA threatens to reveal the true type of the bad firm with a sufficiently high probability, that it chooses not to access to save the fee. But no matter what the CRA threatens to do after an access by a bad firm, consider the following deviation: The bad firm accesses, and the CRA plays exactly the same strategy as if the accessed firm was good. The unsuspecting investor believes it is good. The bad firm gets higher profits than playing no access, and the CRA gets a higher fee than following up on its promised threat. And the “equilibrium” unravels. For the second part of the proof, assume $\mu(A, L) = \mu(A, H)$ and these two objects are determined by Bayes’s rule. Assume $\mu(A, H) > \mu(A, L)$, then the CRA would always say L and would never say H . H would not be on-path.

There are two ways the type might be signaled in equilibrium. Either the CRA helps by its own actions “separate” the types, or it does not. It is possible to get signaling, even without the CRA’s help; the mere existence of the option to access or not, allows for some signaling. Good firms access, for example, and bad firms do not, irrespective of what the CRA does. For some parameters, there is a mixed strategy equilibrium in the one-shot game where the CRA blabs (does the same thing after p_H and p_L), and firms randomize on access and no access making each other indifferent. In this equilibrium, the CRA does not signal by itself. The existence of the option to access or not allows for “some separation” between the types. Roughly speaking, Proposition 2 stated that there is no PBE in which the CRA “helps in separating” the types by its own actions.

5. Infinitely repeated game

The CRA could not credibly commit to reveal the true project type in a one-shot game. This section presents an infinitely repeated game between a long-run CRA and short-run investors and firms. Investors competing to provide funding to firms do not collude, and representing them as short-run

² Firm profits are decreasing in the fee.

players prevents collusion. On the other hand, firms are short-run players because as explained in the introduction, structured finance is specific to the pool of loans at hand. A special-purpose vehicle constructed for a particular pool of loans is better represented as a short-run player. The long-run CRA maximizes its discounted sum of per-period profits. When it chooses a rating, it considers the rating's effect on both its immediate payoff and its discounted future payoffs. Would the rating's effect on future payoffs discipline the CRA when it is sufficiently patient? Would infinitely repeating the game ensure an informationally efficient equilibrium?

5.1. Stage game and setup. Time is discrete and infinite. The stage game is analogous to the one-shot game of 4. In every period t , a short-run firm, firm t , and two short-run investors, investors t_1 and t_2 , play in period t and exit. The CRA is the only long-run strategic player, with discount factor $0 < \delta < 1$. Nature selects a new project every period (i.i.d.). The CRA's payoff is the discounted sum of its per-period payoff. Let $\alpha_t = 0$ if firm t did not access the CRA, and $\alpha_t = 1$ if it did, and let the selected issuance size in period t be b_t . The CRA evaluates a sequence $\{b_t\}_{t=1}^\infty$ as $\sum_{t=1}^\infty \alpha_t \delta^{t-1} \epsilon b_t$

5.2. Public history and strategies. At any time period t , the public history records: whether firm t accessed the CRA or not, the public announcement if access is played, the investors' offers, the chosen investor, and the project realization. A time t public history is: $h_t^p \in [\{NA, (A, H), (A, L)\} \times R^2 \times [0, 1]^2 \times \{1, 2\} \times \{D, P\}]^{t-1}$. Let $h_t^p = \emptyset$. Let H_t^p be the space of all possible time t public histories h_t^p with $H^p = \cup H_t^p$.

Firm t sees the public history from the past and the true project type at period t and randomizes on access and no access, $\sigma_{Firm\ t}: H_t^p \times \{p_H, p_L\} \rightarrow \Delta \{A, NA\}$. The CRA¹ uses only the public history from the past and randomizes on announcements after access by a p_H firm and a p_L firm, $\sigma_{CRA}: H^p \times A \times \{p_H, p_L\} \rightarrow \Delta \{H, L\}$. Investor ti sees the public history from the past and picks an offer for each of the three contingencies of period t , $\sigma_{ti}: H_t^p \times \{NA, (A, H), (A, L)\} \rightarrow \mathcal{R} \times [0, 1]$.

5.3. Perfect Bayesian equilibrium. Since a new project is selected every period i.i.d., incomplete in-

formation about the project's type is transient, and does not carry over from one period to the other.

Definition 4. A PBE is a pair $((\sigma, \mu)$ where $\sigma = ((\sigma_{Firm\ t})_{t=1}^\infty, \sigma_{CRA}, \{\sigma_{i1}\}_{i=1}^\infty, \{\sigma_{i2}\}_{i=1}^\infty)$ is a collection of strategies and $\mu = \{\mu_t\}_{t=1}^\infty$ is a collection of time t beliefs $\mu_t = (\mu(H_t^p \times NA), \mu(H_t^p \times (A, L)), \mu(H_t^p \times (A, H)))$. At each period t , given the public history, fit is three numbers between zero and one that denote investors' beliefs about the probability of facing a PI project, in each of their information sets. Satisfying:

- ◆ Given $\{\mu_t\}_{t=1}^\infty$, $\{(\sigma_{Firm\ t})_{t=1}^\infty$, and $\{\sigma_{i1}\}_{i=1}^\infty, \{\sigma_{i2}\}_{i=1}^\infty$ is optimal at each node the CRA moves on.
- ◆ Given μ_t and H_t^p at each information set, σ_{ii} is optimal given σ_{ij} ; and $\sigma_{Firm\ t}$ is optimal after p_H and after p_L given $\{\sigma_{CRA}, \sigma_{i1}, \sigma_{i2}\}$.
- ◆ μ_t is derived from Bayes's rule when possible, with no restrictions where it does not apply.

5.4. Informationally efficient equilibrium. The efficiency concept in the infinitely repeated game is similar to, but slightly more intricate than, section 4.3's. Efficiency requires that in every period, on the equilibrium path investors are informed of the period's project type. The equilibrium path is where investors' beliefs are formed by Bayes's rule².

Definition 5. A PBE is informationally efficient if, at every time period t on the equilibrium path, investors t_1 and t_2 have correct beliefs about the type of the time t project.

As the discount factor goes to 1, the continuation payoff's weight in the CRA's total payoff increases, and the CRA cares more about the future effect of its rating. Is there an informationally efficient PBE as the discount factor approaches 1?

Proposition 3. $\forall \delta: 0 \leq \delta < 1$, there is no informationally-efficient equilibrium in the infinitely repeated game.

Proof. An informationally efficient equilibrium reveals the true project type at every non-measure zero contingency of every period t . As in Proposition 2, the true project type of period t is revealed in three scenarios. Two cases require the bad firm to access and have its true type revealed. This would never happen in equilibrium, since the worst that

¹ A player's strategy might possibly depend on anything he knows of what has occurred in the past, be it private or public. But here, I impose the standard restriction that players only use public events of the past. The players' strategies still depend on what they see privately today, but it can depend on the past only through what is publicly known.

² To further clarify what the PBE definition entails, assume that at period 1 investors 11 and 12 experienced a measure zero contingency. A PBE puts no restrictions on their period 1 beliefs in that contingency. In period 2, time 2 investors should not have seen a measure zero contingency in the public history, but that is irrelevant for determining their period 2 beliefs about the period 2 project. If they do see the actions they expect in period 2, then their beliefs will again be determined by Bayes's rule. For the intricacies of PBE, check Fudenberg and Tirole (1991).

could happen after no access is to be thought bad, but even then it saves the fee. The only possible case then is: The bad firm plays no access, the good firm accesses, and the CRA threatens to make an access decision by the bad firm unprofitable. I will find a profitable deviation and conclude the proof. Let the selected issuance size after access, when the investors know it is a good firm, be:

$$b_H^* = 1 - \left[\frac{p_H}{(1-p_H)(1-\epsilon)R-1} \right]^{-p_H}. \quad \text{An informationally}$$

efficient equilibrium has informationally efficient continuations after every possible contingency. So, the CRA's continuation payoff after every possible contingency arising at period t is constant at $\delta(1-\eta)b_H^*$. The deviation is that the bad firm accesses, the CRA says what it says after a good firm accesses, and the investors are fooled. The bad firm gets higher profits because it is thought to be good. The CRA gets a higher payoff today, and the continuation payoff is constant, netting it a higher total payoff.

Proposition 3 shows that no informationally efficient equilibrium exists. But what if the "inefficiency" is only a finite time phenomenon? Can there be a point in time after which play will be efficient?

Corollary 1. *There is no asymptotically informationally efficient equilibrium $\exists t^*$, an $H_{t^*}^p$, and a PBE (σ, μ) s.t. the continuation equilibrium strategies and beliefs after the "public" subgame starting with $H_{t^*}^p$, $(\sigma|H_{t^*}^p, \mu|H_{t^*}^p)$, is an informationally efficient equilibrium in the game starting with the "public node" $H_{t^*}^p$.*

Proof. If there was then, by properly redefining the strategies and beliefs, an informationally efficient equilibrium of the original game would have existed.

Corollary 1 shows that "inefficiency" is not a transient phenomenon, resolved in finite time. Because of discounting, this phenomenon is not an "inefficiency" that can only happen "at infinity". Hence, I conclude that "inefficiency" is a persistent and recurring phenomenon in every PBE.

5.5. A simple no access and a truthful-recursive equilibrium. I already showed the nonexistence of an informationally efficient equilibrium, this section constructs two equilibria. A fairly general trend in applications, following Abreu et al.'s (1990) approach, tends to characterize the equilibrium payoff set. Here, the only strategic player is the CRA, and characterizing its payoff set is not of much help in understanding investors' equilibrium outcomes, which are the main objects of interest. This section constructs a simple PBE and a truthful recursive

PBE¹, where the CRA fulfills its role for only a finite number of periods². The simple equilibrium delivers the CRA's worst punishment. Although the "inefficiency" in every PBE is persistent and recurrent, can I use the simple equilibrium to construct a truthful equilibrium that separates the project types for at least a finite number of periods?

Lemma 4. *Under Condition 2 and Condition 3, the following is a simple PBE: Every good, and bad firm t does not access the CRA. In every period, after access by a p_L and a p_H firm, the CRA puts probability $\frac{1}{2}$ on H . Investors t_1 , and t_2 act optimally, keeping their priors after no access, and believing it is a bad firm after access.*

Proof. No firm wants to access, pay the fee, and be considered bad; hence no firm would deviate. A CRA deviation has no effect on its payoff (given that investors always believe it is a bad firm that accessed, irrespective of the announcement). Investors act optimally given their beliefs, which are derived by Bayes's rule when possible.

5.5.1. Best truthful recursive equilibrium with z forgiveness. In the simple equilibrium, the CRA does not signal anything by its own actions, and investors retain their priors on-path. I now construct an equilibrium, where the CRA is accessed only by the good firm, and is truthful up to the point where it announces H , and the project defaults. Players then either forgive or move to the no-access equilibrium forever. The details follow:

$$\text{Let } b_H^* = 1 - \left[\frac{p_H}{((1-\epsilon)\bar{R}-1)(1-p_H)} \right]^{-1-p_H}, \text{ and}$$

$$b_L^* = 1 - \left[\frac{p_L}{((1-\epsilon)\bar{R}-1)(1-p_L)} \right]^{-1-p_L}$$

be the issuance sizes bought after access, when investors are sure it is the good and the bad project respectively. Consider an automaton with two phases:

- ◆ Truthful recursive phase: The CRA tells the truth. Good firms access the CRA, and bad firms do not. Investors act optimally given their beliefs. After access and H , they believe the firm is good; after access and L or no access, they believe it is bad.

¹ The PBEs constructed in this section are also sequential equilibria. It is well known that the sequential equilibrium concept does not put many restrictions beyond PBE, in this kind of signaling game.

² The fact that play converges to a punishment state, where investors are not informed, is crucial in understanding the nondesirability of these equilibria. The CRA is informed of the type, and the absorption into the "state of uninformiveness" is highly undesirable. The focus generally is on where equilibrium play eventually converges and not on the play in the (finite) beginning of the game.

- ◆ Punishment no-access phase: The CRA randomizes $\frac{1}{2}$ on H and $\frac{1}{2}$ on L . Neither firm accesses. Investors act optimally given their beliefs. After access, they believe the firm is bad; after no access, they retain their prior.

All players start at the truthful phase, and remain there, until the following happens: the CRA is accessed, it announces H , and the project defaults. After this contingency, with probability z , the CRA is forgiven, and players remain in the truthful recursive phase; but with $1 - z$, players move to the punishment no-access phase, and remain there forever. Call this a z -automaton. For every δ , we are most interested in the equilibrium with the highest possible z , since that would make the transition to the punishment no-access phase least probable, and would raise the probability that investors are informed.

Proposition 4. *Under Condition 2 and Condition 3, $\exists 0 < \underline{\delta} < 1$ s.t. $\forall \delta : \underline{\delta} \leq \delta, \exists \bar{z}(\delta) : 0 \leq \bar{z} < 1$ s.t. $\bar{z} = \max\{0 \leq z < 1 : \text{the } z\text{-automaton is a PBE}\}$.*

Proof. The no-access punishment phase is the no-access equilibrium of Lemma 4. In the truthful separating phase, Condition 3 guarantees that a good firm would rather access, and pay the fee to signal its true type, than not access and be considered a bad firm. A bad firm prefers no access (where it is known to be bad); since access also reveals it is a bad firm, and it pays the fee. Investors act optimally given their beliefs, and the beliefs are consistent with the strategies proposed. What remains is to check the CRA’s incentive constraints.

Since the CRA is not accessed in the punishment no-access phase, its value in that phase is zero. Assume the value of the truthful recursive phase is v .

Promise keeping: $v = \frac{(1-\eta)(1-\delta)\epsilon b_H^*}{1-\delta[(1-(1-\eta)(1-z)p_H]}$. The

incentive constraints after access by p_H and by p_L follow: $(1-\delta)\epsilon(b_H^* - b_L^*) \geq \delta p_H(1-z)v$ and $(1-\delta)\epsilon(b_H^* - b_L^*) \leq \delta p_L(1-z)v$. Substituting in promise keeping:

$$\begin{aligned} \frac{\delta(1-z)p_H(1-\eta)}{1-\delta[(1-(1-\eta)(1-z)p_H]} &\leq \frac{b_H^* - b_L^*}{b_H^*} \leq \\ &\leq \frac{\delta(1-z)p_L(1-\eta)}{1-\delta[(1-(1-\eta)(1-z)p_H]} \end{aligned}$$

Note, that b_H^* and b_L^* depend neither on δ nor on z . The LHS and RHS are strictly decreasing in z , and strictly increasing in δ . Let $\underline{\delta}$ be the unique solution

to $\frac{\delta p_L(1-\eta)}{1-\delta[(1-(1-\eta)p_H]} = \frac{b_H^* - b_L^*}{b_H^*}$. If $\delta < \underline{\delta}$ the incentive

constraint after p_L would not hold, no matter what z is. The CRA would never say that a bad firm is bad. So let $\underline{\delta} \leq \delta$, and define $\bar{z}(\delta)$ to be the unique z that

solves: $\frac{b_H^* - b_L^*}{b_H^*} = \frac{\delta(1-z)p_L(1-\eta)}{1-\delta[(1-(1-\eta)(1-z)p_H]}$. \bar{z} is

where the incentive constraint after a bad firm access binds. At this \bar{z} , the other constraint necessarily holds. And at every $z > \bar{z}$, the incentive constraint after the bad firm access is violated. This concludes the proof.

I note that $\bar{z}(\delta)$ is increasing in δ . This is intuitive because the more patient a CRA is, the less punishment is needed.

6. Infinitely repeated game with verification

This section allows for rating verification. After investment decisions are made, with some fixed preset probability $0 < p < 1$, the true project type is publicly revealed. Random verification is interpreted in one of two ways. In the first, investors have full confidence in the verification results. With some preset probability p , either investors are committed to verify the rating, or analysts verify the rating process and publicly disseminate their opinions of the CRA’s work. In the second, investors or analysts are committed to always verify, and p is investors’ confidence in the verification method¹. With $1-p$, investors view the verification as irrelevant, and disregard its conclusions. Structured finance has a lower p because these products are more “complex” and/or more “opaque” (less of the information underlying them is public). On the other hand, the facts that much of the information underlying the bond is public and that bonds are generally “simple” products allow for a higher probability of verification, and/or stronger investor confidence in the verification method and results.

6.1. Verification probability, cutoff discount factor and efficiency. I show here that for every strictly positive verification probability, there exists a cutoff CRA discount factor, above which an informationally efficient equilibrium exists. Verification allows for punishments tailored to CRA deviations that actually occur. No matter how small (but strictly positive) the verification is, when the CRA is patient enough, the punishment is potent enough to deter any deviation.

¹ There is no lack of impending-doom prophets, no matter what the economic environment is (Dr. Doom is one example). p specifies investors’ confidence in Dr. Doom’s forecast.

Proposition 5. *In the infinitely repeated game with verification, under Condition 2 and Condition 3, for any verification probability $0 < p \leq 1$, $\exists \underline{\delta}(p): 0 < \underline{\delta} < 1$ s.t. $\forall \underline{\delta} \leq \delta$, there exists an informationally-efficient equilibrium.*

Proof. Omitted details of this proof follow the proof of Proposition 4. The following is a PBE (for the δ s specified below). There are two phases: In both the truthful phase and the punishment phase, players act and believe as in the equilibrium of Proposition 4. The transition function is that all players start at the truthful phase, and stay there until a lie is detected; once that occurs, they switch to the punishment phase, and remain there forever. In the punishment phase, no player wants to deviate because of (a straightforward adaptation of) Lemma 4. The no-access equilibrium is still an equilibrium in the game with verification. In the truthful phase, I focus on the CRA's incentive constraints and derive the cutoff δ . The CRA's value in the punishment phase is v . Assume the value of the truthful phase is v . Promise keeping: $v = (1 - \eta)eb_H^*$. The CRA's incentive constraint after p_H access is $(1 - \delta)eb_L^* + \delta(1 - p)v \leq (1 - \delta)eb_H^* + \delta v$; and it always holds. If it says H , the CRA gets eb_L^* today and tomorrow it always gets v , whether there was verification or not. If it lies and says L , it gets eb_L^* today and only gets v if there was no verification. The CRA's incentive constraint after p_L access is $(1 - \delta)eb_H^* + \delta(1 - p)v \leq (1 - p)eb_L^* + \delta v$. If it says L , the CRA gets eb_L^* today and always gets v tomorrow. If it lies and says H , it gets a higher payment eb_H^* today, but only gets v if there was no verification. Solving the p_L access constraint gives us the cut off delta $\underline{\delta}$. Substitute promise-keeping in that constraint to get $\underline{\delta} = \frac{b_H^* - b_L^*}{b_H^* [p(1 - \eta) + 1] - b_L^*}$. If $\underline{\delta} < \delta$, the constraints hold and the CRA does not deviate.

6.2. Discount factor, cutoff verification probability and efficiency. In subsection 6.1, I fixed the verification probability and derived the CRA's cutoff patience level above which, an informationally efficient equilibrium exists. It is of more interest to know for every fixed discount factor, the required verification level above which an informationally efficient equilibrium exists. Compared with the discount factor, the verification probability is more easily controlled and changed.

Proposition 6. *In the infinitely repeated game with verification, $\exists \underline{\delta}: \underline{\delta} > 0$ s.t.*

- ◆ If $\delta < \underline{\delta}$, then $\forall p: 0 \leq p \leq 1$ there is no informationally efficient equilibrium.

- ◆ If $\underline{\delta} < \delta$, then $\exists p(\delta): 0 \leq p(\delta) \leq 1$ s.t. If $p \leq p(\delta)$, there is no informationally efficient equilibrium.

Proof. An informationally efficient equilibrium puts restrictions on the continuation payoffs allowable. Let $v = (1 - \eta)eb_H^*$. If the CRA acts truthfully, then the continuation payoff should be v . Let $\underline{\delta} = \frac{b_H^* - b_L^*}{b_H^* - b_L^* + (1 - \eta)b_H^*} > 0$. If $\underline{\delta} < \delta$, then $(1 - \delta)eb_L^* + \delta v < (1 - \delta)eb_H^* \leq (1 - \delta)eb_H^* + \delta(1 - p)v \forall p$. Hence, when $\underline{\delta} < \delta$, the CRA is not truthful after a p_L firm access. At those low deltas, the short term gain dominates. If $\underline{\delta} < \delta$, then let $p(\delta) = \frac{(1 - \delta)(b_H^* - b_L^*)}{\delta(1 - \eta)b_H^*}$. If $p < p(\delta)$, consider the constraint after a p_L access with continuation payoffs that make it most easily satisfied (the RHS and LHS continuations at maximum and minimum consistent with efficiency): $(1 - \delta)eb_H^* + \delta(1 - p)v < (1 - \delta)eb_L^* + \delta v$. If $p < p(\delta)$, the constraint would be violated. If $p(\delta) \leq p$, the constraint holds, and the strategies and beliefs of before where good firms access, bad firms do not, beliefs on-path are determined by Bayes's rule, off-path the firm is considered bad, and the CRA punishment of forever no access, if caught lying, form an informationally efficient equilibrium.

The intuition for Proposition 6 is straightforward. When $\delta < \underline{\delta}$, the CRA (future) punishment has no bite. No matter how much it is verified, it will never act truthfully. If the CRA cares enough about the future $\underline{\delta} \leq \delta$, then for each δ there is a cutoff verification probability, above which verification is likely enough, and the CRA acts truthfully. As δ increases and the CRA becomes more patient, the cutoff verification required to sustain an informationally efficient equilibrium decreases.

Conclusion

This paper showed that rating verifiability is crucial for mitigating CRA conflicts of interests. Models with insufficient rating verifiability, as in the structured finance market, show that the CRA does not fully disclose its information about the project's probability of default and therefore does not fulfill its role, even after endogenizing the reputation function in an infinitely repeated game. Models with enough verification, as in the bond market, show that there is an equilibrium where the CRA fulfills its role as the information asymmetry mitigator.

When the rating is verifiable, the interaction between the probability of verification and the CRA patience level is analyzed. I show two main results: The first looks for required patience level given the verification

probability: for any nonzero verification probability, if the CRA is patient enough with respect to that probability, then it fulfills its role. The second looks for the required verification level for any patience level: when the CRA is not patient enough, there is no efficient equilibrium, even with certain verification. If the CRA is patient enough, verification has to be likely enough to guarantee efficiency. This paper forms the background of a recommendation for investors to carefully verify structured finance products. Frequent and convincing verification, coupled with a potent punishment when verification shows the CRA deviated, forms an

essential part of providing enough incentives for the CRA to reveal its information truthfully.

Acknowledgement

The author is grateful to Christopher Phelan, Jan Werner, Joe Haubrich, David Rahman, Aldo Rustichini, Futoshi Narita, and Filippo Occhino for advice and suggestions, seminar participants at Minnesota, Cleveland FED, Ryerson University, Monash University, University of Sydney, Midwest Economic Conference, Kent State University, and American University of Beirut for helpful comments.

References

1. Abreu, D., Pearce, D., and Stacchetti, E. (1990). Towards a theory of discounted repeated games with imperfect monitoring, *Econometrica*, 58, pp. 1041-1063.
2. Ashcraft, A. and Schuermann, T. (2008). Understanding the securitization of sub-prime mortgage credit, Working paper, Federal Reserve Bank of New York.
3. Benmelech, E. and Dlugosz, J. (2009). The credit rating crisis, NBER Macroeconomics Annual, pp. 161-207.
4. BIS (2005). The role of ratings in structured finance: Issues and implications, *Bank of International Settlements*, 23.
5. Bolton, P., Freixas, X., and Shapiro, J. (2012). The credit ratings game, *Journal of Finance*, LXVII (1).
6. Cantor, R. and Packer, F. (1994). The credit rating industry. FRBNY Quarterly Review.
7. Carlson, M. and Hale, G. (2006). Rating agencies and sovereign debt rollover, *Topics in Macroeconomics*, 6 (2) Article 8.
8. Coval, J., Jurek, J., and Stafford, E. (2009). The economics of structured finance, *Journal of Economic Perspectives*, 23, pp. 3-26.
9. Covitz, D. and Harrison, P. (2003). Testing conflicts of interest at bond rating agencies with market anticipation: Evidence that reputation incentives dominate, Working Paper.
10. Elamin, M. (2010). Can reputation ensure efficiency in the structured finance market? Working paper, University of Minnesota.
11. Fudenberg, D. and Tirole, J. (1991). *Game Theory*, MIT Press, Cambridge, MA.
12. Griffin, J. and Tang, D. (2010). Did subjectivity play a role in cdo credit ratings? Working paper, UTAustin.
13. Kendall, L. and Fishman, M. (2000). *A Primer on Securitization*, MIT Press, Cambridge.
14. Kuhner, C. (2001). Financial rating agencies: Are they credible? Insights into the reporting agencies in times of enhanced systemic risk, *Schmalenbach Business Review*, 53, pp. 2-26.
15. Lizzeri, A. (1999). Information revelation and certification intermediaries. *RAND Journal of Economics*, 30 (2), pp. 214-231.
16. Mailath, G. and Samuelson, L. (2006). *Repeated Games and Reputations: Long-Run Relationships*, Oxford University Press, Oxford.
17. Mathis, J., McAndrews, J., and Rochet, J.-C. (2009). Rating the raters: Are reputation concerns powerful enough to discipline the rating agencies? *Journal of Monetary Economics*, 56, pp. 657-674.
18. Rablen, M. (2013). Divergence in credit ratings, *Finance Research Letters*, 10, pp. 12-16.
19. Sangiorgi, F. and Spatt, C. (2012). Opacity, credit rating, shopping and bias, Working paper, Stockholm School of Economics and Carnegie Mellon.
20. SEC (2008). Summary report of issue identified in the commission staff's examination of select credit rating agencies. Available at: <http://www.sec.gov/news/studies/2008/craexamination070808.pdf>.
21. Skreta, V. and Veldkamp, L. (2009). Ratings shopping and asset complexity, *Journal of Monetary Economics*, 56 (5), pp. 678-695.
22. Smith, R. and Walter, I. (2001). Rating Agencies: Is There an Agency Issue? NYU Working Paper No. FIN-01-003. Available at SSRN: <http://ssrn.com/abstract=1294460>.
23. Titman, S. and Tsyplakov, S. (2010). Originator performance, cmts structures and the risk of commercial mortgages, *American Real Estate and Urban Economics (AREUEA) Annual Meetings conference paper*.
24. White, L. (2010). Markets: The credit rating agencies, *Journal of Economic Perspectives*, 24, pp. 211-226.

Appendix

Proof of Proposition 1. Let $q = \eta p_L + (1 - \eta) p_H$ be the posterior on default when the prior is retained. Let σ_1 with offer (R_1, b_1) and σ_2 with offer (R_2, b_2) and σ_{firm} form a PBE. Obviously, given $((R_1, b_1), (R_2, b_2))$, if the firm is called upon to move, it will pick i whenever: $(\bar{R} - R_1) b_i > (\bar{R} - R_j) b_j$ where $i, j \in \{1, 2\}$. For ease, given $((R_1, b_1), (R_2, b_2))$, let i 's expected utility in case his offer is accepted be: $E_i = q \log(1 - b_i) + (1 - q) \log(1 + (R_i - 1) b_i)$ and the profits to the firm when i 's offer is accepted and the project pays be: $P_i = (\bar{R} - R_i) b_i$. This proof comes in steps. I show that both equilibrium offers give the firm the same profit and investors the same expected utility. Then, that the expected utility is zero. Finally, I show the offers give the firm the maximum profit, subject to the two constraints.

Step 1. Assume $P_1 \neq P_2$, w.l.o.g. let $P_1 < P_2$. I will find a profitable deviation offer for 2, (R'_2, b'_2) , that gives 2 expected utility E'_2 and the firm profit P'_2 .

- ◆ If $b_2 = 0$, then $P_1 < P_2 = 0$. 2 is selected for sure, but nothing is invested. Condition 1 guarantees that the derivative of E_2 (when $R = \bar{R}$) at $b_2 = 0$ is strictly positive. The value is zero at $b_2 = 0$ hence $\exists \tau > 0$ s.t. with $b'_2 = \tau$ and $R'_2 = \bar{R}$, the new offer (R'_2, b'_2) is still chosen for sure, since $P_i < P'_2 = 0$ and 2 gets a higher expected utility $E_2 = 0 < E'_2$.
- ◆ If $b_2 \neq 0$, then $\exists r > 0$ s.t. with $b'_2 = b_2$ and $R'_2 = R_2 + r$ making a new offer (R'_2, b'_2) requiring a bit more return than before, 2's offer is still chosen for sure, since $P_1 < P_2$ and 2 gets a higher expected utility $E_2 < E'_2$.

Step 2. With $P_1 = P_2$, assume $E_1 \neq E_2$, w.l.o.g. $E_1 < E_2$. Then 1 is selected with probability $\frac{1}{2}$ and gets a lower expected utility than 2. 2's offer is the deviation.

Step 3. With $P_1 = P_2$, assume $E_1 = E_2 \neq 0$. First, if $E_1 = E_2 < 0$, then 1 is better off eating his endowment (i.e., to offer (P_1, b_1) with $R_1 < P_2$). Second, if $E_1 = E_2 > 0$ then $\exists \tau > 0$ s.t. with $b'_2 = b_2 \neq 0$ and $R'_2 = R_2 - \tau$ making a new offer (R'_2, b'_2) requiring just a bit less return, 2 would be chosen for sure, $P_1 = P_2 < P'_2$, consequently he gets a higher expected utility $E_2 < E'_2$, hence 2 would deviate to this new offer.

Step 4. We know $E_1 = E_2 = 0$ and $P_1 = P_2$. Now let $P = \text{Max}_{(R,b)} (\bar{R} - R)b$ s.t. $0 \leq b < 1$ and $E = 0$. Assume $P_1 = P_2 < P$.

- ◆ If $P_2 < 0$ and $E_2 = 0$, then 1 offers $R'_1 = \bar{R}$ and any $b'_1 > 0$ that gives him an $E_1 > 0$, and he is selected for sure, since $P_1 = 0 > P_2$.
- ◆ If $P_2 \geq 0$, then there is always an alternative offer for 1 that would give the firm higher profits and 1 a higher expected utility. 1 either offers (R_1, b_1) with $R_1 > R_2$ and $b_1 > b_2$ or 1 offers (R_1, b_1) with $R_1 < R_2$ and $b_1 < b_2$. The following graph details the better offer by 1 that gives him a positive expected utility and gives the firm higher profits.

The colored regions represent the contours of the expected utility of the investors; as the numbers on the contours show, the darker the color the lower the investors' expected utility. The two bold black contours represent the contours of the firm's profits, and the arrows show the direction of increasing profits. The grey circle is the offer that solves the problem P1.

If 2 offers the black box (which gives expected utility zero), then by decreasing both the issuance size and the required rate of return, 1 can offer the contract represented by the red circle. At the red-circle offer, the firm gets a higher profit, and 1 gets a higher expected utility than zero (the red circle will be the deviation offer if the black box is offered). Similar reasoning governs the deviation offer if the lower intersection of the black contour intersects the zero expected utility contour. The deviation offer would increase the issuance size and the required return. This concludes the proof.




Fig. 1. Required return graphic representation

Proof of Lemma 1. This proof has an existence part and a uniqueness part. Drop constraint $0 \leq b \leq 1$ and solve by Lagrange to get:

$$b^* = 1 - \left[\frac{q}{(1-q)(\bar{R}-1)} \right]^{1-q} \text{ and } R^* = \frac{1}{q} \left[(1-q)\bar{R} - 1 \right]^{\frac{1-b^*}{b^*}}. \text{ Condition 1 implies that } 1 < \frac{1}{1-q} < \bar{R}, \text{ and we have that } b^* < 1.$$

And since Condition 1 implies that $\frac{q}{((1-q)\bar{R}-1)} < 1$, we have that $b^* > 0$.

Existence. Let $\underline{b} = \frac{b^*}{2}$ and $\bar{b} = \frac{1+b^*}{2}$, we have $0 < b < b^* < \bar{b} < 1$. $[\underline{b}, \bar{b}]$ is compact. Now the constraint $q \log [1 - b] + (1 - q) \log [1 + (R - 1) b] = 0$ defines a function from b to R in the following manner: for $b \in [\underline{b}, \bar{b}]$, $R(b) = \frac{1 - (1 - b)^{\frac{1}{1-q}}}{b(1 - b)^{\frac{q}{1-q}}}$. The function $R(\cdot)$ is continuous on $[\underline{b}, \bar{b}]$, hence it maps the compact set to a compact set. So, $[\underline{b}, \bar{b}] \times R[\underline{b}, \bar{b}]$, is compact. The objective function is continuous on a compact set, Heine-Borel's theorem ensures existence.

Uniqueness. The constraint gives the following required return for any $b \in [\underline{b}, \bar{b}]$, $R(b) = \frac{1 - (1 - b)^{\frac{1}{1-q}}}{b(1 - b)^{\frac{q}{1-q}}}$. Substituting

in the objective function, we have: $g(b) = Rb + 1 - b - \frac{1}{(1 - b)^{\frac{q}{1-q}}}$. The objective function is strictly concave on the domain,

the constraint set is convex, and hence there is a unique b^* that solves the problem. The function $R(b) = \frac{1 - (1 - b)^{\frac{1}{1-q}}}{b(1 - b)^{\frac{q}{1-q}}}$ defines

the unique R^* that corresponds to the unique b^* .