



EUROPEAN SUMMER SYMPOSIUM IN FINANCIAL MARKETS

**Generously hosted by
Study Center Gerzensee**

Monday 13 July-Friday 24 July 2009

Ratings Shopping and Asset Complexity: A Theory of Ratings Inflation

Vasili Skreta (New York University) and *Laura Veldkamp (New York
University)

We are grateful to the following institutions for their financial and organizational support: Study Center Gerzensee and the Swiss National bank.

The views expressed in this paper are those of the author(s) and not those of the funding organization(s) or of CEPR, which takes no institutional policy positions.

RATINGS SHOPPING AND ASSET COMPLEXITY: A THEORY OF RATINGS INFLATION

Vasiliki Skreta and Laura Veldkamp *

New York University, Stern School of Business

February 2009

Abstract

Many identify inflated credit ratings as one contributor to the recent financial market turmoil. We develop an equilibrium model of the market for ratings and use it to examine possible origins of and cures for ratings inflation. In the model, asset issuers can shop for ratings – observe multiple ratings and disclose only the most favorable – before auctioning their assets. When assets are simple, agencies' ratings are similar and the incentive to ratings shop is low. When assets are sufficiently complex, ratings differ enough that an incentive to shop emerges. Thus, an increase in the complexity of recently-issued securities could create a systematic bias in disclosed ratings, despite the fact that each ratings agency produces an unbiased estimate of the asset's true quality. Increasing competition among agencies would only worsen this problem. Switching to an investor-initiated ratings system alleviates the bias, but could collapse the market for information.

Life can only be understood backwards; but it must be lived forwards.
Soren Kierkegaard (1813 - 1855)

*vskreta@stern.nyu.edu and lveldkam@stern.nyu.edu, 44 West Fourth St., Suite 7-180, New York, NY 10012. This paper was prepared for the Carnegie-Rochester conference series. Many thanks to David Backus, Ignacio Esponda, Valerie Bencivenga, Dimitri Vayanos and Lawrence White for useful discussions. We also thank participants in the Stern Micro Lunch for their helpful suggestions.

Most market observers attribute the recent credit crunch to a confluence of factors: excess leverage, underestimation of risk, opacity, lax screening by mortgage originators, improperly estimated correlation between bundled assets, market-distorting regulations, a rise in the popularity of new asset classes whose risks were difficult to evaluate, as well as credit rating agency conflicts of interest.¹ This paper investigates the misrating of structured credit products, widely cited as one contributor to the crisis. Our main objective is to critically examine two arguments about why ratings problems arose and to show how combining the two could produce a ratings bias that imperfectly informed investors would not anticipate.

One argument focuses on asset issuers who shop for the highest ratings. The New York Times explains:²

The banks pay only if [the ratings agency] delivers the desired rating . . . If Moody's and a client bank don't see eye-to-eye, the bank can either tweak the numbers or try its luck with a competitor like S&P, a process known as *ratings shopping*.

While the issuer-initiated ratings system has been around since the 1970s,³ ratings bias only recently emerged as a concern. To argue that it took 30 years to detect the bias is to suggest that financial-market participants learn unrealistically slowly. This raises the question: Is it possible that ratings shopping previously had either no or a small effect and that something about the credit market changed to amplify its effect?

A second argument about why many assets were systematically misrated attributes the problem to the increasing complexity of assets. As Mark Adelson testified before Congress,⁴

The complexity of a typical securitization is far above that of traditional bonds. It is above the level at which the creation of the methodology can rely solely on mathematical manipulations. 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.

However, the credit market crisis was not generated by independent ratings errors. Only systematic

¹See, for instance, page 1 of the Memorandum for the President from the President's Working Group on financial-markets dated March 13, 2008.

²Quote from *New York Times Magazine*, "Triple-A-Failure," April 27, 2008. Other articles making similar arguments include "Why Credit-rating Agencies Blew It: Mystery Solved," available from <http://robertreich.blogspot.com/2007/10/they-mystery-of-why-credit-rating.html>; "Stopping the Subprime Crisis" *New York Times*, July 25, 2007; "When It Goes Wrong" *The Economist*, September 20, 2007, and "Credit and Blame" *The Economist*, September 6, 2007.

³The credit-rating industry switched from "the investor-pays" to "the issuer-pays" model in the early 1970s. This is believed to have happened due to the widespread availability of photocopier technology. See, for instance, White (2007).

⁴Adelson: Director of structured finance research at Nomura Securities. Testimony before the Committee on Financial Services, U.S. House of Representatives, September 27, 2007. On January 26, 2008, *The New York Times* quoted Moody's CEO, saying "In hindsight, it is pretty clear that there was a failure in some key assumptions that were supporting our analytics and our models." He said that one reason for the failure was that the "information quality" given to Moody's, "both the completeness and veracity, was deteriorating." See also page 10 of the *Summary Report of Issues Identified in the Commission Staff's Examinations of Select credit-rating Agencies*, United States Securities and Exchange Commission, July 8, 2008.

upward ratings would produce a widespread rise in the prices of credit products. This raises the question: Is it possible that more dispersion in ratings can translate into higher ratings on average?

We show that the combination of an increase in asset complexity and the ability of asset issuers to shop for ratings can produce ratings inflation, even if each rating agency produces an unbiased rating. We do not argue that the complexity of any given asset increased. Rather, the composition of assets being sold changed so that the more complex type of asset, the structured financial products- particularly those that were mortgage-backed- became more prevalent. For example, while under \$10 billion in structured finance collateralized debt obligations (CDO's) were distributed in 2000, nearly \$200 billion were issued in 2006 (Hu 2007).

The intuition behind our results is as follows: Each rating agency issues an unbiased forecast of an asset's value. However, if the announced rating is the maximum of all realized ratings, it will be a biased signal of the asset's true quality. The more ratings differ, the stronger are issuers' incentives to selectively disclose (shop for) ratings. For simple assets, agencies issue nearly identical forecasts. Asset issuers then disclose all ratings because more information reduces investors' uncertainty and increases the price they are willing to pay for the asset. For complex assets, ratings may differ, creating an incentive to shop for the best rating. There is a threshold level of asset complexity such that once this threshold is crossed, shopping becomes optimal and ratings inflation emerges. Furthermore, the link between asset complexity and ratings shopping can work in both directions. An issuer who shops for ratings might want to issue an even more complex asset, to get a broader menu of ratings to choose from. This, in turn, makes shopping even more valuable.

Biased ratings affect securities prices if investors are unaware of the bias. If investors do not know that the complexity of assets has changed, then based on past data, they would rationally expect ratings to be unbiased, until they observe a sufficient amount of data to detect the bias.

It is possible that there was no ratings bias and that this episode is simply a bad draw, producing correlated errors in ratings. Because of the persistence in asset returns and the short sample history of many of the new credit-derivative products, empirically establishing that their ratings were biased will be possible only in the far future. However, the pattern of ratings suggests a relationship between asset complexity and over-optimistic ratings. Mason and Rosner (2007) document that complex CDOs have significantly higher default rates than simple corporate bonds with identical ratings.⁵ Similarly, mortgage-backed securities whose underlying credit risk, correlation risk and pre-payment risk are notoriously difficult to assess, experienced more widespread downgrades than did assets based on other collateral types.⁶ Lastly,

⁵ Alan Greenspan acknowledged the greater complexity of CDOs in his May 2005 testimony, "[T]he credit risk profile of CDO tranches poses challenges to even the most sophisticated market participants." He cautioned investors "not to rely solely on rating-agency assessments of credit risk."

⁶ "Other collateral types that began to be securitized well after mortgages are far less complex. The first non-mortgage securitization was equipment leases, followed by credit cards and auto loans, and, more recently, home equity, lease finance, manufactured housing, student loans, and synthetic structures. All of those types of collateral illustrate tranching structures that are measurably simpler than those for RMBS." (Mason and Rosner 2007)

even if the primary cause of misrating was bad luck, a failure to address potential sources of bias inherent in the structure of the ratings industry could generate future problems.

Section 1 models a market for ratings with a few salient, realistic features. In reality, the two largest agencies, Moody's and S&P, account for 80 percent of market share. When a structured credit product is issued, the issuer typically proposes a structure to an agency and asks it for a "shadow rating." This rating is private information between the agency and the issuer, unless the issuer pays the agency to make the rating official and publicize it. In the model, an asset issuer can purchase and make public one or two signals about the payoff of an asset. We call these signals "ratings." After choosing the number of ratings to observe and which ones to make public, the issuer holds an auction for his assets. After each investor submits a menu of price-quantity pairs, the asset issuer sets the highest market-clearing price for his asset, and all investors pay that price per share.

As a benchmark, Section 2 solves this model with mandatory disclosure of all observed ratings. Section 3 solves the model with voluntary disclosure. Our main results are in section 4. On average, obtaining and disclosing a rating benefits an asset issuer, because information makes asset payoffs less uncertain, and investors pay more for less-risky assets. Interestingly, the value of obtaining an additional rating is non-monotonic in asset complexity. If assets become more complex and harder to rate, then using additional ratings to ratings shop becomes more profitable. If assets become very complex, then ratings are so noisy that investors largely ignore them. This makes asset issuers less willing to pay the issuance fee for a rating that will have little impact on the price of their asset. Furthermore, if an asset issuer can choose to make his asset more complex, then knowing that he will shop for ratings can make more asset complexity desirable.

Section 5 uses our model to evaluate the effect of recently-proposed reforms. First, we consider whether allowing only investor-initiated ratings is viable or desirable. We modify the model to allow any investor to observe an asset's rating at a cost. To create the potential for ratings bias, we subject a fraction of investors to investment-grade securities regulation, which restricts them to only purchasing assets whose ratings exceed a minimum threshold. Informed ratings expand their investment-possibility set. However, investors do not shop for ratings because they optimally use all available information in forming their bids. Thus, if the source of bias is not ratings agencies themselves, but is, instead, ratings shopping, then investor-initiated ratings are likely to be less biased. The downsides to investor-initiated ratings are that investors can free-ride on others' information and that the ratings market can easily collapse. If investors cannot overcome these problems, even biased ratings may be better than no information at all.

Another proposed reform is to allow more agencies to compete in providing ratings. For the case of issuer-initiated ratings, we show that having more rating agencies would exacerbate the problem of bias because there would be more ratings from which to choose. In the investor-initiated ratings model, increasing the number of agencies is inconsequential. Finally, we consider the logistical problems with mandatory disclosure

laws.

A controversial assumption we maintain throughout is that agencies produce unbiased ratings. Appendix B explores a version of our model that relaxes this assumption. We do not deny that rating agencies may report biased ratings in an attempt to increase their business. Rather, our point is that even if rating agency conflicts of interest are resolved, bias could continue to plague ratings.

Our contribution vis-a-vis the existing literature As a theoretical contribution, our paper builds on three distinct literatures. First, we add an interaction between ratings and equilibrium asset prices to the literature on ratings agencies. Other papers that model ratings agencies are quite distinct from our work. Faure-Grimaud, Peyrache and Quesada (2007) identify circumstances under which the optimal renegotiation-proof contract between the rating agency and the firm results in the firm owning its rating. In their setup, the rating reveals the asset value perfectly, and the price of the asset is exogenous. Farhi, Lerner and Tirole (2008) focus on other aspects of ratings, such as their transparency and coarseness. In Damiano, Li and Suen (2008), Bolton, Freixas and Shapiro (2008) and Becker and Milbourn (2008), a rating agency prefers to inflate its clients' ratings, but has some reputation cost of reporting a value different from the asset's true, exogenous price. Sangiorgi, Sokobin and Spatt (2008) develop an equilibrium explanation for credit-rating "notching" and the connection between rating shopping and the "winner's curse." The aforementioned papers investigate the equilibrium level of bias. In contrast, our model's rating agencies report the truth. We show that even if ratings agencies produce unbiased ratings, bias in disclosed ratings can still exist.

Second, we extend the literature on information in asset markets by modeling disclosure incentives.⁷ In existing models with endogenous information, investors can acquire unbiased signals. Our model explores how asset issuers choose to both acquire and disclose signals, and the resulting signal bias.

Third, we augment the existing literature on sellers who provide information about their goods by considering how much information to provide.⁸ In Shavell (1994), either the seller or a buyer can acquire information. As in our paper, Shavell studies equilibrium information acquisition under voluntary and mandatory disclosure regimes. But because Shavell's buyers know that sellers may be hiding information, any equilibrium without full disclosure unravels. In Jovanovic (1982), the fact that disclosure is costly prevents unraveling. In all these papers, there is a single signal to reveal or not, which precludes the possibility of shopping for ratings.

More broadly, our findings highlight the role that institutions, rules and market structure play in an industry that produces information. A central question in the mechanism-design literature is: Which institutions

⁷For an overview of the role of exogenous information in asset markets, see Brunnermeier (2001). For more recent work on this topic, see Banerjee (2007). Seminal works on endogenous information include Grossman and Stiglitz (1980) and Hellwig (1980). More recent work on investor-initiated information acquisition includes Barlevy and Veronesi (2000), Bullard, Evans and Honkapohja (2005) and Peress (2004).

⁸Seminal contributions in this literature include Grossman (1981) and Milgrom (1981). Benabou and Laroque (1992), Morgan and Stocken (2003) and, more recently, Bolton, Freixas and Shapiro (2007) analyze the conflict of interest between the buyer and seller in this environment.

are most desirable when information is asymmetric or dispersed? This paper asks the reverse question: What information do agents choose to observe or disclose in a given institution and market structure? As the recent crisis highlights, understanding the information provision is as important as understanding the institutions. When information production runs amok, large economic fluctuations can result.

1 A Model of a Market for Ratings

This is a static model of an *asset issuer*, who has many units of an asset to sell and a continuum of *investors* who want to buy those assets. The asset's value is unknown to investors. Credit rating agencies produce and sell information. The total supply of the asset is fixed and determined by the issuer. The market-clearing price is determined through a uniform-price auction in which the sum of the bidders' bidding schedules determines the aggregate demand.⁹ Investors choose their bidding functions so as to maximize their utility, subject to the information that they have. What information they have depends on whether the rating is purchased by the issuer (in which case, the issuer makes it public) or by the investors. Below, we investigate how each of these different arrangements about who purchases the information affects the quality and amount of information available to market participants.

We now describe a model in which issuers pay for ratings. Later in the paper, we analyze a version of this model in which investors pay for ratings.

Assets There are two assets: The 'safe' asset offers riskless return r , and the risky asset pays u , which is normally distributed $u \sim N(\bar{u}, \sigma_u^2)$. The price of the riskless asset is 1. The price of the risky asset is p , which is endogenous.

Investors A continuum of ex-ante identical investors has utility

$$U = -e^{-\rho(m_i r + q u)}, \quad (1)$$

where ρ is the coefficient of absolute risk aversion and q_i and m_i are the number of risky and riskless asset shares investor i ends up with. Each agent is endowed with m_i^0 units of the riskless asset, but can borrow and lend that asset freely at the riskless rate r . Hence, each investor's budget constraint is

$$m_i + p q_i = m_i^0. \quad (2)$$

⁹This auction is similar to the limit economy in Reny and Perry (2006), but we incorporate investor risk-aversion and budget constraints.

The Auction The price of the risky asset is determined in an auction. Each investor submits a bidding function that specifies the maximum amount that he is willing to pay for q units of the risky asset as a function of his information. These bid functions determine the aggregate demand. The auctioneer specifies a market-clearing price p that equates aggregate demand and supply, and each trader pays this price for each unit purchased (uniform-price auction).

We now state the bid function and verify that it constitutes an equilibrium. Bids depend on each investor's information set \mathcal{I}_i , which includes information inferred from b being the price paid per unit.

$$b(q|\mathcal{I}_i) = \frac{E(u|\mathcal{I}_i) - q\rho V(u|\mathcal{I}_i)}{r}, \quad (3)$$

where $E(u|\mathcal{I}_i)$ and $V(u|\mathcal{I}_i)$ are the mean and variance of the risky asset's return, conditional on the investor's information. The price paid per unit is exogenous from each investor's perspective because he is infinitesimal compared to the rest of the market, implying that the price he faces is determined by other investor's bid functions, together with the aggregate supply. Therefore, a realized price b reveals information about others' bids, which, in turn, is partially informative about what they know.

Each bidder is infinitesimal, which implies that he takes the market-clearing price as given. Thus, the bidding function (3) is the inverse demand function of a trader who seeks to maximize (1) subject to (2), taking p as given. It can be easily verified that the objective function of this constrained maximization problem is concave in q , so that the first-order condition describes the optimal portfolio:

$$q_i = \frac{1}{\rho} V[u|\mathcal{I}_i]^{-1} (E[u|\mathcal{I}_i] - pr). \quad (4)$$

Because (3) is an inverse of (4), it is a best response given everyone else's bid function.

When issuers solicit a rating, they either disclose the rating to all investors or keep it private so that no investor observes it. Either way, investors have symmetric information \mathcal{I} . Integrating over the asset demand (4) and equating aggregate demand with the asset supply, delivers the equilibrium price

$$p = \frac{1}{r} (E[u|\mathcal{I}] - \rho \text{Var}[u|\mathcal{I}]). \quad (5)$$

Rating Agencies Credit-rating agencies produce ratings, which are noisy, unbiased signals about the risky asset payoff u . We consider two rating agencies, because this is the simplest setting in which we can illustrate our results.

We assume that a shadow rating θ is an unbiased signal about the payoff - that is, $\theta \sim N(u, \sigma_\theta^2)$, - produced at marginal cost χ . The issuer can choose to keep the rating private or to make it public information.¹⁰ All

¹⁰If an issuer does not publish the shadow rating, rating agencies are not withholding the rating because the exact structured product they rated is rarely issued. Rather, the rating and the asset structure are negotiated. Another agency may tell the

rating agencies produce the same service. Since there is no quantity choice, firms compete in a Bertrand way and set price equal to marginal cost: χ for shadow ratings and $\chi + \chi$ for publicly-issued ratings.

Assets vary with respect to the degree that they can be reliably rated, an attribute we call asset complexity.

Definition 1 *A more complex asset is one with more noise in its ratings: It has a higher σ_θ^2 .*

Asset Issuer The issuer knows the true payoff[®] of the risky asset and is endowed with x shares of it. His objective is to maximize expected profit. We first consider the case in which the issuer initiates the rating. In this model, the issuer's expected profit is the price times quantity of the asset sold, minus the cost of observed and disclosed ratings:

$$\Pi = px - s\chi - s\chi, \quad (6)$$

where s is the number of shadow ratings observed, including the ones eventually disclosed, and s is the number of publicly disclosed ratings.

Model Timing *Stage 1: The issuer's ratings-acquisition decision.* The issuer decides whether or not to acquire a shadow rating. If he decides to do so, he visits one of the two rating agencies. Upon acquiring and observing the first shadow rating, he decides whether or not to acquire another shadow rating.

Stage 2: The issuer's ratings-disclosure decision. Depending on the number of acquired shadow ratings, the issuer decides whether to disclose (publish) no, one, or both ratings.

Stage 3: Price Determination. An auction determines the market-clearing price.

The asset issuer's decisions with voluntary disclosure are summarized in Figure 1. After these stage-1 decisions are made, the asset auction takes place.

2 The Benchmark: Mandatory Disclosure of Shadow Ratings

With mandatory disclosure, the issuer must only decide is whether to acquire zero, one or two ratings. He does so sequentially: First, he decides whether to acquire one rating, and then, upon observing the realization of the first rating, he decides whether it is worthwhile to obtain a second.

The expected price of an unrated asset is¹¹

$$p_0 \equiv \frac{1}{r}(\bar{v} - \rho\sigma_u^2x). \quad (7)$$

issuer that structuring the security with slightly more low-risk assets will earn it the sought-after rating. Then, this slightly modified security would be issued. See Mason and Rosner (2007) for a detailed description of this process.

¹¹Some fixed-income securities are issued without ratings. Such unrated bonds are classified as junk bonds. Typically, these are small-net-worth assets.

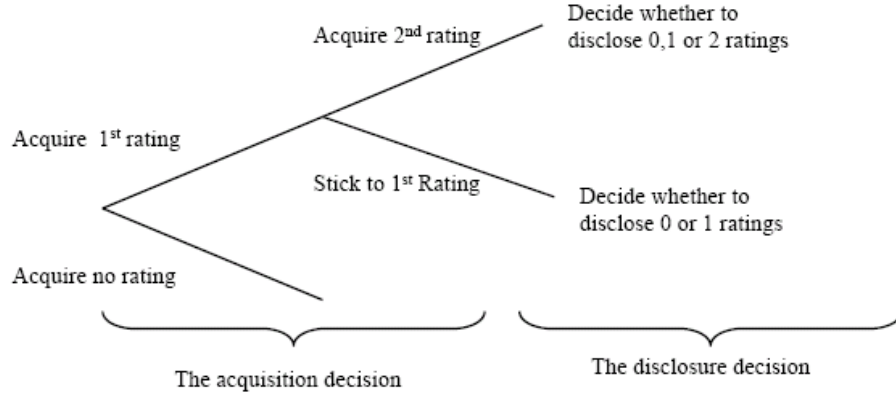


Figure 1: An asset issuer's decision tree.

If the issuer initiates a rating, Bayes' law dictates that the expected value of the asset is $E[u|\theta] = (\sigma_u^{-2}\bar{u} + \sigma_\theta^{-2}\theta)/(\sigma_u^{-2} + \sigma_\theta^{-2})$, and the conditional variance of the asset will be $V[u|\theta] = 1/(\sigma_u^{-2} + \sigma_\theta^{-2})$. Since the issuer decides to acquire the rating before he knows its outcome, he considers the expected price

$$\bar{p}_1 \equiv \frac{1}{r}(\bar{u} - \rho V[u|\theta]x) = \frac{1}{r} \left(\bar{u} - \frac{\rho x}{\sigma_u^{-2} + \sigma_\theta^{-2}} \right). \quad (8)$$

Thus, the difference in issuer profits from buying information is

$$\bar{\pi}_{s=1}^M - \bar{\pi}_{s=0}^M = \frac{\rho x^2}{r} \frac{1}{\sigma_\theta^2 \sigma_u^{-4} + \sigma_u^{-2}} - \chi - \bar{\chi}, \quad (9)$$

where $\bar{\pi}_{s=0}^M$ ($\bar{\pi}_{s=1}^M$) stands for the issuer's expected profits from acquiring no (respectively, one) rating under mandatory disclosure. The issuer chooses to purchase a rating if (9) is non-negative.

The asset's expected price with two ratings is

$$\bar{p}_2 = \frac{1}{r} \left(\bar{u} - \frac{\rho x}{\sigma_u^{-2} + 2\sigma_\theta^{-2}} \right). \quad (10)$$

The issuer chooses to obtain two shadow ratings instead of one if

$$\bar{\pi}_{s=2}^M - \bar{\pi}_{s=1}^M = \frac{\rho x^2}{r} \frac{1}{\sigma_\theta^2 \sigma_u^{-4} + 3\sigma_u^{-2} + 2\sigma_\theta^{-2}} - \chi - \bar{\chi} > 0. \quad (11)$$

Note, that if (9) is positive, (11) will be, too. Hence, if the issuer profits from acquiring one shadow rating, he also profits from a second.

Comparative Statics with Mandatory Disclosure Information resolves risk. Therefore, the more

risk-averse investors are, the more they value information. The value of a rating is also increasing in the quantity of the asset offered because the benefit of information can be applied to every share. Finally, the value of a rating is non-monotonic in asset complexity (σ_θ). The value of the first rating is always decreasing in complexity because a more-complex asset is harder to rate, and the resulting rating is less precise and, thus, less valuable ((9) is decreasing in σ_θ^2). The value of the second rating could also be lower for the same reason, or it could be higher because having less information from the first rating increases the marginal value of additional information: (9) is increasing in asset complexity if the variance of the asset's returns is high, relative to the complexity of the asset ($2\sigma_u^4 > \sigma_\theta^2$), and is decreasing otherwise.¹²

Ultimately, a model with mandatory disclosure of ratings by asset issuers and truthful reporting by rating agencies cannot explain the bias in ratings. We introduced it because it illustrates the mechanics of the solution. We now explore the more realistic voluntary-disclosure case to understand where ratings bias might come from.

3 Solving the Voluntary Disclosure Model

The key trade-off an asset issuer faces is the following: Withholding the most negative ratings makes the asset appear more valuable to investors, while publicizing more ratings makes the asset less risky. In other words, disclosure lowers the conditional variance of the risky-asset payoff u , while ratings shopping increases its conditional mean. Both effects increase the price investors are willing to pay and, thus, the issuer's profit. We investigate which circumstances favor ratings shopping, as well as the resulting bias in the asset price.

For this discussion to be meaningful, we need to ensure that ratings bias is not irrelevant. In particular, sophisticated investors should not be able to infer the expected bias, subtract that expected bias from the announced rating, and neutralize its effects through their actions.

Assumption 2 *Investors do not correct for ratings selection bias: For every announced rating, they believe that $\theta \sim N(u, \sigma_\theta^2)$. For unrated assets, they believe that $u \sim N(\bar{u}, \sigma_u^2)$.*

This assumption implies that investors cannot make correct inferences from the rating agencies' actions (the number of ratings they choose to disclose). Players being unable to form the correct mapping about the informativeness of other people's actions is an important feature of the equilibrium notion of Eyster and Rabin (2005), and of Esponda (2008) and is also present in the analysis of DeMarzo, Vayanos and Zwiebel (2003), among others.

The assumption that investors believed ratings to be unbiased is consistent with our main argument that much of the bias was a recent phenomenon. Suppose that the situation we model were repeated, and asset

¹² $\frac{\partial(\Pi_{s=2}^M - \Pi_{s=1}^M)}{\partial \sigma_\theta} = -\frac{2\sigma_\theta \sigma_u^{-4} - 4\sigma_\theta^{-1}}{(\sigma_\theta^2 \sigma_u^{-4} + 3\sigma_u^{-2} + 2\sigma_\theta^{-2})^2}$, is positive if $\frac{4}{\sigma_\theta} > 2\sigma_\theta \sigma_u^{-4}$ or $2 > \sigma_\theta^2 \sigma_u^{-4}$ or $2 > \frac{\sigma_\theta^2}{\sigma_u^4}$ or $2(\sigma_u^2)^2 > \sigma_\theta^2$.

complexity suddenly changed. If investors did not observe the change, but knew that regime changes in complexity were possible, they would infer asset complexity and, thus, ratings bias from the past history of ratings and asset outcomes. Since the historical data came mainly from simple assets, investors would initially believe that assets are simple, that no ratings shopping is taking place, and that ratings are unbiased. Even after assets became more complex, this belief would persist until they observed a sufficiently long series of ratings and payoffs from the complex assets. Thus, with an unexpected change in asset characteristics, even rational investors would not have initially detected ratings bias.¹³

3.1 The disclosure decision

To solve the model, we start with the last decision and work backwards. We begin by considering an issuer who has already chosen how many ratings to acquire and is deciding how many to disclose. If the asset issuer chooses not to solicit any ratings, then the asset price is the same whether disclosure is voluntary or mandatory because there is no rating to disclose. But when the asset issuer solicits one or two shadow ratings, he faces the following choices.

Disclosure with two shadow ratings We first investigate the disclosure decision of an issuer who has acquired two shadow ratings. We call the higher rating $\hat{\theta}$, and the lower $\underline{\theta}$, so that $\hat{\theta} > \underline{\theta}$. We want to identify under which conditions the issuer will disclose none, one or both ratings. Since the asset issuer is always more inclined to announce a higher rating, disclosing one rating means disclosing $\hat{\theta}$.

We first compare the alternatives of disclosing one versus no ratings. If the issuer announces no ratings, the conditional mean and variance of the asset payoff are the unconditional mean and variance, \bar{x} and σ_u^2 . Therefore, the price of the asset is the same as in (7).

If the issuer announces rating $\hat{\theta}$, the price will be

$$p_1(\hat{\theta}) = \frac{1}{r} \left(\frac{\sigma_u^{-2}\bar{x} + \sigma_{\hat{\theta}}^{-2}\hat{\theta} - \rho x}{\sigma_u^{-2} + \sigma_{\hat{\theta}}^{-2}} \right). \quad (12)$$

Let $\pi_{D=d}$ stand for the issuer's profit from disclosing d ratings. Then, the additional profit gained from disclosing information is

$$\pi_{D=1}(\hat{\theta}) - \pi_{D=0} = [p_1(\hat{\theta}) - p_0]x - \chi, = \frac{x}{r} \frac{\hat{\theta} - \bar{x} + \rho x \sigma_u^2}{1 + \sigma_u^{-2} \sigma_{\hat{\theta}}^2} - \chi. \quad (13)$$

Having purchased at least one shadow rating, the asset issuer discloses that rating if (13) is positive. Let a

¹³If investors are sophisticated and can perfectly account for how the issuers' disclosure rules vary with asset complexity, we anticipate that unraveling would occur similar to that in Shavell (1994), where all undisclosed information is treated as if it is bad news. Thus, all ratings would be disclosed.

be the value of $\bar{\theta}$ that causes (13) to be zero:

$$a = \frac{r(1 + \sigma_u^{-2}\sigma_{\bar{\theta}}^2)}{x}\chi + \bar{u} - \rho x \sigma_u^2. \quad (14)$$

Since (13) is monotonically increasing in $\bar{\theta}$, the asset issuer discloses at least one rating if $\bar{\theta} \geq a$. In summary, when disclosure of ratings is voluntary, the issuer discloses one versus no ratings when the rating obtained is high enough.

Consider, next, the choice between disclosing one or two ratings. If the issuer discloses both $\bar{\theta}$ and $\underline{\theta}$, the asset price will be

$$p_2(\bar{\theta}, \underline{\theta}) = \frac{1}{r} \frac{\sigma_u^{-2}\bar{u} + \sigma_{\bar{\theta}}^{-2}(\bar{\theta} + \underline{\theta}) - \rho x}{\sigma_u^{-2} + 2\sigma_{\bar{\theta}}^{-2}}. \quad (15)$$

The issuer prefers to disclose both ratings if

$$\begin{aligned} \Pi_{D=2}(\bar{\theta}, \underline{\theta}) - \Pi_{D=1}(\bar{\theta}) &= [p_2(\bar{\theta}, \underline{\theta}) - p_1(\bar{\theta})]x - \chi, \\ &= \frac{x}{r} \frac{(\sigma_{\bar{\theta}}^{-2}\sigma_u^{-2}(\underline{\theta} - \bar{u}) + \sigma_{\bar{\theta}}^{-4}(\underline{\theta} - \bar{\theta}) + \rho x \sigma_{\bar{\theta}}^{-2})}{(\sigma_u^{-2} + 2\sigma_{\bar{\theta}}^{-2})(\sigma_u^{-2} + \sigma_{\bar{\theta}}^{-2})} - \chi > 0. \end{aligned} \quad (16)$$

Since (16) is monotonically increasing in $\underline{\theta}$, the issuer discloses both ratings if $\underline{\theta} \geq b(\bar{\theta})$, where $b(\bar{\theta})$ is the value of $\underline{\theta}$ that equates (16) to zero when the highest rating is $\bar{\theta}$:

$$b(\bar{\theta}) = \frac{1}{(\sigma_{\bar{\theta}}^{-2}\sigma_u^{-2} + \sigma_{\bar{\theta}}^{-4})} \left(\chi \frac{r(\sigma_u^{-2} + 2\sigma_{\bar{\theta}}^{-2})(\sigma_u^{-2} + \sigma_{\bar{\theta}}^{-2})}{x} + \sigma_{\bar{\theta}}^{-2}\sigma_u^{-2}\bar{u} + \sigma_{\bar{\theta}}^{-4}\bar{\theta} - \rho x \sigma_{\bar{\theta}}^{-2} \right). \quad (17)$$

In other words, the issuer prefers to show- that is, disclose only the highest of the two ratings- when $\underline{\theta}$ is sufficiently lower than $\bar{\theta}$, ($\underline{\theta} < b(\bar{\theta})$).

Finally, the issuer will disclose no ratings if no ratings are preferable to one rating and to two ratings. Both these conditions are satisfied if $\bar{\theta} < a$.¹⁴

In summary, the disclosure decision for an issuer that has acquired two shadow ratings is,

$$\begin{aligned} \text{Disclose both ratings if } & \bar{\theta} \geq a \text{ and } \underline{\theta} \geq b(\bar{\theta}). \\ \text{Disclose highest rating if } & \bar{\theta} \geq a \text{ and } \underline{\theta} < b(\bar{\theta}). \\ \text{Disclose no ratings if } & \bar{\theta} < a. \end{aligned}$$

¹⁴Zero ratings are preferred to two ratings when $\Pi_{D=0} - \Pi_{D=2}(\bar{\theta}, \underline{\theta}) > 0$, which happens if $\left[\frac{1}{r}(\bar{u} - \rho \sigma_u^2 x) - \frac{1}{r} \frac{\sigma_u^{-2}\bar{u} + \sigma_{\bar{\theta}}^{-2}(\bar{\theta} + \underline{\theta}) - \rho x}{\sigma_u^{-2} + 2\sigma_{\bar{\theta}}^{-2}} \right] x + 2\chi > 0$. When $\bar{\theta} \leq a$, the highest possible sum for $\bar{\theta} + \underline{\theta} = 2 \left(\frac{r(1 + \sigma_u^{-2}\sigma_{\bar{\theta}}^2)}{x}\chi + \bar{u} - \rho x \sigma_u^2 \right)$ (this is because, by definition, $\underline{\theta} \leq \bar{\theta}$). Substituting this sum in the above inequality, one can see that it is always satisfied. The details can be found in Technical Appendix: Computation Details for the Disclosure Decision 2 versus 0. This observation allows us to conclude that the issuer will disclose no ratings when $\bar{\theta} < a$.

Disclosure with one shadow rating Suppose that the asset issuer has acquired only one shadow rating. With a slight abuse of notation, we call that rating $\hat{\theta}$. The issuer prefers to disclose if (13) is positive (when $\hat{\theta} \geq a$) and does not disclose otherwise.

3.2 The acquisition decision

Now that we understand the issuer's disclosure decisions, we move back one node in the decision tree to study ratings' acquisition. The issuer makes two decisions sequentially. First, he decides whether to acquire the first shadow rating, and then he decides whether to acquire the second one. Again, we work backwards: We start with the decision of the issuer who already has one shadow rating and considers whether to obtain a second. We call θ_1 the first shadow rating observed and θ_2 the second.

The decision to acquire the second rating The decision depends on whether the first rating is high enough to disclose ($\theta_1 \geq a$).

Case 1: The first rating was high ($\theta_1 > a$). If the second draw is low relative to the first draw ($\theta_2 < b(\theta_1)$), then the issuer discloses only the first rating θ_1 . If the second draw is sufficiently high to disclose and not so high that it makes the first rating no longer worthwhile to disclose ($b(\theta_1) < \theta_2 < \theta^*$), the issuer discloses both ratings. If the second rating is sufficiently higher than the first rating, the issuer discloses only the second rating. Thus, if the first rating is such that $\theta_1 > a$, the expected price contains three terms corresponding to these three possibilities:

$$\hat{p}^{(2)}(\theta_1) = F(b(\theta_1))p_1(\theta_1) + \int_{b(\theta_1)}^{\theta^*} p_2(\theta_1, \theta_2)f(\theta_2)d\theta_2 + \int_{\theta^*}^{\infty} p_1(\theta_2)f(\theta_2)d\theta_2, \quad (18)$$

where $p_1(\theta_1)$ and $p_2(\theta_1, \theta_2)$ are defined in (12) and (15), and where θ^* is the value of θ_2 such that $\theta_1 = b(\theta_2)$, implying that $\theta^* = b^{-1}(\theta_1)$. Because the asset issuer knows the true value of the asset, the probability density $f(\theta_2)$ is conditional on u and is independent of the realization of θ_1 .

Then, whenever $\theta_1 > a$ and the issuer sticks with one rating, his profit is

$$\pi^{(1)}(\theta_1) = p_1(\theta_1)x - \chi - \chi,$$

whereas if he acquires a second rating, his expected profit conditional on θ_1 is

$$\pi^{(2)}(\theta_1) = \hat{p}^{(2)}(\theta_1)x - 2(\chi + \chi), \quad (19)$$

where $\hat{p}^{(2)}(\theta_1)$ is given by (18). The expected benefit from acquiring the second rating depends on the difference in the asset price between disclosing two ratings versus disclosing one, and on the difference in

the price between disclosing θ_2 alone and disclosing θ_1 alone. Each difference is weighted by the probability that θ_2 takes on a value that makes the associated disclosure optimal. Hence, when $\theta_1 > a$, we have

$$\begin{aligned}
p^{(2)}(\theta_1) - p^{(1)}(\theta_1) &= \int_{b(\theta_1)}^{\theta_1^{-1}(\theta_1)} [x(p_2(\theta_1, \theta_2) - p_1(\theta_1)) - \chi] f(\theta_2) d\theta_2 \\
&\quad + \int_{b^{-1}(\theta_1)}^{\infty} x(p_1(\theta_2) - p_1(\theta_1)) f(\theta_2) d\theta_2 - \chi, \\
&= \int_{b(\theta_1)}^{\theta_1^{-1}(\theta_1)} \left[\frac{x}{r} \frac{(\sigma_\theta^{-2} \sigma_u^{-2} (\theta_2 - \theta_1) + \sigma_\theta^{-4} (\theta_2 - \theta_1) + \rho x \sigma_\theta^{-2})}{(\sigma_u^{-2} + 2\sigma_\theta^{-2}) (\sigma_u^{-2} + \sigma_\theta^{-2})} - \chi \right] f(\theta_2) d\theta_2 \\
&\quad + \int_{b^{-1}(\theta_1)}^{\infty} \left(\frac{x}{r} \left(\frac{\sigma_u^{-2} \theta_2 + \sigma_\theta^{-2} (\theta_2 - \theta_1) - \rho x}{\sigma_u^{-2} + \sigma_\theta^{-2}} \right) \right) f(\theta_2) d\theta_2 - \chi. \tag{20}
\end{aligned}$$

where the second line uses (12) and (15) to substitute out p_1 and p_2 .

Case 2: The first rating was low ($\theta_1 < a$). If the second rating is also low ($\theta_2 < a$), then the issuer will disclose no ratings. If the second rating is moderately high, it is possible that the issuer discloses both ratings, even though the first was too low to disclose on its own. If the second rating is high ($\theta_2 > b^{-1}(\theta_1)$), the issuer discloses only the second rating. If $\theta_1 < a$, the expected price contains three terms that correspond to these three possibilities,

$$p^{(2)}(\theta_1) = F(a)p_0 + \int_a^{b^{-1}(\theta_1)} p_2(\theta_1, \theta_2) f(\theta_2) d\theta_2 + \int_{b^{-1}(\theta_1)}^{\infty} p_1(\theta_2) f(\theta_2) d\theta_2, \tag{21}$$

where p_0 is given by (7).

If $\theta_1 < a$ and the issuer does not obtain a second rating, then he discloses no rating and earns profit $p^{(1)}(\theta_1) = p_0 x$. If he acquires a second rating, his expected profit conditional on θ_1 is

$$p^{(2)}(\theta_1) = p^{(2)}(\theta_1 | \theta_1 < a) x - 2(\chi + \chi). \tag{22}$$

Thus, if $\theta_1 < a$, the expected benefit from acquiring the second rating is

$$\begin{aligned}
&p^{(2)}(\theta_1) - p^{(1)}(\theta_1) \\
&= \int_a^{b^{-1}(\theta_1)} [(p_2(\theta_1, \theta_2) - p_0) x - 2\chi] f(\theta_2) d\theta_2 + \int_{b^{-1}(\theta_1)}^{\infty} [(p_1(\theta_2) - p_0) x - \chi] f(\theta_2) d\theta_2 - \chi \\
&= \int_a^{b^{-1}(\theta_1)} \left(\frac{x((\theta_1 + \theta_2) - 2\theta_1 + 2\rho x \sigma_u^2)}{r(2 + \sigma_\theta^2 \sigma_u^{-2})} - 2\chi \right) f(\theta_2) d\theta_2 + \int_{b^{-1}(\theta_1)}^{\infty} \left(\frac{x(\theta_2 - \theta_1 + \rho x \sigma_u^2)}{r(1 + \sigma_u^{-2} \sigma_\theta^2)} - \chi \right) f(\theta_2) d\theta_2 - \chi. \tag{23}
\end{aligned}$$

Let E_2 denote the region of realizations of the first rating where the differences in (20) and in (23) are

positive. For this region, the issuer will choose to obtain a second draw.

The decision to acquire the first rating When the issuer makes this decision, he has only his prior information about the possible realizations of the ratings. He compares the expected profit from acquiring no rating, $\pi_i^{(0)} = \frac{x}{r}(1 - \rho\sigma_u^2 x)$, with the expected profit from acquiring the first rating, $\pi_i^{(1)}$. This profit is calculated anticipating four possible future acquisition and disclosure decisions: The first rating may be too low to disclose but may prompt the investor to acquire a second rating; the first rating may be high enough to disclose and still prompt the acquisition of a second rating; the first rating may be too low to disclose, and may still deter the acquisition of a second rating; and, finally, the first rating may be high enough to disclose and may deter the acquisition of a second rating:

$$\begin{aligned} \pi_i^{(1)} = & \int_{[-\infty, a] \cap \Theta_2} \pi_i^{(2)}(\theta_1) dF(\theta_1) + \int_{[a, \infty] \cap \Theta_2} \pi_i^{(2)}(\theta_1) dF(\theta_1) \\ & + \int_{[-\infty, a] \cap \Theta_2^c} p_0 x dF(\theta_1) + \int_{[a, \infty] \cap \Theta_2^c} (p_1(\theta_1)x - \chi - \chi) dF(\theta_1), \end{aligned} \quad (24)$$

where p_0 is given by (7), p_1 is given by (12), and the formula for expected profits with two ratings $\pi_i^{(2)}(\cdot)$ is in (19) when $\theta_1 > a$ and in (22) when $\theta_1 < a$.

4 Main Results: Asset Complexity and Ratings Bias

So far, we have analyzed the issuer's incentives to acquire shadow ratings under two regimes: the mandatory- and voluntary- disclosure regimes. For each of these regimes, we identified cases in which the issuer prefers to obtain none, one, or two shadow ratings. In each case, choices depend on the characteristics of the asset to be rated. One of those characteristics is the asset's complexity.

This section considers the interaction between greater asset complexity and the ratings shopping that creates ratings bias. There are two pieces to this interaction. First, we consider how an exogenous change in asset complexity affects the incentive to shop for ratings. Second, we show that ratings shopping can create an incentive to structure more-complex securities.

4.1 How complexity affects the incentive to shop for ratings

Our argument is that ratings shopping arose as the nature of credit products and their market changed. For this to be a plausible explanation, we need to show that the benefit of ratings shopping in (16) is increasing in some parameter that was trending up at the time. This leads us to ask the following comparative statics question: Given two shadow ratings, what happens to the incentive to publicly disclose these ratings as an asset becomes more complex?

The effect of complexity turns out to be non-monotonic. For either very low or very high asset complexity ($\sigma_\theta^2 \rightarrow 0$ or $\sigma_\theta^2 \rightarrow \infty$), ratings shopping always takes place. When $\sigma_\theta^2 \rightarrow \infty$ in (16), $\mathbb{E}_{D=1}(\hat{\theta}) - \mathbb{E}_{D=2}(\hat{\theta}, \underline{\theta})$ clearly converges to the cost of disclosure χ . When $\sigma_\theta^2 \rightarrow 0$, the same result follows from the fact that $\hat{\theta} - \underline{\theta}$ tends to zero. Since χ is positive, disclosing one rating is preferred to disclosing both. The intuition is that when asset complexity is small, ratings are precise. The extent to which publicizing a second rating reduces the risk of investing in the asset is too small to be worth the cost. When asset complexity is high, ratings become uninformative. Since investors know this, issuing multiple ratings has little price impact and is, again, not worth the cost.

Numerical example We consider an asset issuer who has observed two shadow ratings and is deciding how many of these to disclose. The net benefit of disclosing a second rating is given by (16). Since this depends on the realized ratings, we first take an expectation over these realizations. For a normal variable with mean \bar{u} and variance σ_θ^2 , the expected highest- and lowest-order statistics out of a sample of two are $\frac{\sigma_\theta}{\pi} + \bar{u}$ and $-\frac{\sigma_\theta}{\pi} + \bar{u}$ (Kotz, Balakrishnan and Johnson 2000).¹⁵

Figure 2 shows a case in which the expected profit of disclosing the second rating is lower, higher and then lower again compared to the expected profit from disclosing one rating. The vertical lines represent the two roots $\{\}$ levels of σ_θ that make the issuer indifferent between disclosing one or two ratings. Ratings bias would arise in the first or third regions, where the issuer chooses to disclose only the higher of the two observed ratings.

Not every set of parameters will generate this pattern. If both roots are not positive real numbers, then disclosing one or no ratings could be optimal. But in that case, it would never be optimal to disclose two ratings, for any level of asset complexity.




Figure 2: The average asset issuer's expected value of disclosing zero, one or two ratings, for various levels of asset complexity (σ_θ^2). Parameter values: $x = 2.5$, $r = 1.03$, $\sigma_u^2 = .5$, $\bar{u} = u = 10$, $\chi = 0.3$ and $\rho = 1$.

¹⁵The expectation of $\bar{\theta}$, is the expectation of the highest order statistic out of a sample of two. From Kotz et al., we know that this expectation for a standard normal distribution is given by $\frac{1}{\pi}$. Since, here, ratings are distributed by $N(\bar{u}, \sigma_\theta^2)$, this expectation is given by $\frac{\sigma_\theta}{\pi} + \bar{u}$. This follows from the fact that the expectation is linear in the mean and standard deviation.

This example was constructed for the average signal realizations, in order to keep it simple. As asset complexity increases, there will be a distribution of signal outcomes. Some will prompt asset issuers to disclose them as ratings, others not. Thus, for a given set of parameter values, we can calculate the probability of ratings shopping. Then, instead of a discrete change from no bias to bias, there is a continuous change from a low to a high probability of ratings shopping. The next set of results takes into account the random nature of the realized ratings.

4.2 How complexity affects the demand for shadow ratings

Just as complexity has a non-monotonic effect on disclosure, it also has a non-monotonic effect on ratings acquisition. The reason that its effect is non-monotonic can be seen by considering its limiting cases again. When complexity approaches zero, each rating is perfectly precise. Therefore, there is no benefit and only a cost to acquiring a second rating. When complexity approaches infinity, ratings are uninformative. Again, there is no benefit and only a cost to acquiring either a first or second rating. In between these two extremes, there can be a value to acquiring more than one rating in order to shop for the best one. But if that is the case, then rising complexity must cause the net benefit of a second rating to rise and then fall.

A key effect of changing complexity is that it changes the distribution of $\hat{\theta}$ and $\underline{\theta}$. To incorporate this effect, we do a change of variables and express the expected benefit of acquiring a second rating ((20) and (23)) in terms of standard normal random variables. The partial derivative with respect to σ_θ^2 reveals that the effect of complexity on ratings acquisition is non-monotonic.¹⁶




Figure 3: The asset issuer's expected value of obtaining zero, one or two shadow ratings for various levels of asset complexity (σ_θ^2). Parameter values are identical to those in Figure 1, except that we take expectations over θ_1 and θ_2 , instead of fixing their values. In addition, $\chi = 0.1$.

To see what effects arise for non-extreme parameter values, we return to our numerical example. Using the same parameters as in Figure 2, Figure 3 plots the asset issuer's expected profit from obtaining zero,

¹⁶Details available upon request.

one or two shadow ratings. As asset complexity rises, the number of shadow ratings acquired starts at one, rises to two, falls back to one, and eventually becomes zero.

4.3 How complexity affects ratings bias

The average percentage bias in ratings appears in Figure 4. The concurrent increase in asset complexity and ratings bias is the amalgamation of many different effects. First, it incorporates the disclosure decision in Figure 3. As complexity rises, a firm with two average ratings discloses one, then two, and then one rating again. But Figure 3 obscures the fact that the issuer will not always draw average ratings. Therefore, there will always be some probability of ratings shopping whenever ratings are acquired. The second effect is that as complexity rises, two ratings become farther apart, on average. Thus, the ratings bias generated by ratings shopping grows. The third effect is the change in ratings acquisition. As complexity grows from a low level, more issuers get a second rating, enabling them to choose the highest one to disclose.




Figure 4: The percentage bias in ratings for various levels of asset complexity (σ_θ^2). Parameter values are identical to those in Figure 2. Ratings bias is the average of all disclosed ratings, minus the true mean of the distribution from which the ratings are drawn. It is expressed as a percentage deviation from the true mean.

When complexity becomes very large, bias plummets. This corresponds to the level of complexity where issuers no longer want to acquire any shadow rating because ratings contain too little information to be worth their cost. One might wonder why bias does not drop at $\sigma_\theta = 1.4$, where the value of acquiring no ratings surpasses the value of one rating in Figure 3. This is because the solid line is the value of acquiring one, and only one, rating. But choosing to acquire a first rating also gives the issuer the option value to acquire a second rating. This total value of acquiring a first shadow rating is surpassed by the value of observing no ratings at $\sigma_\theta = 2$. When no ratings are observed, ratings bias disappears.

4.4 When do asset issuers prefer complex assets?

Consider an asset issuer who can choose whether to structure his asset as a simple or as a complex security. Both the simple and the complex security have the same payoff distribution, but the variance of the ratings σ_θ^2 is higher for the more complex asset. The issuer chooses his asset's complexity before observing its shadow ratings.

Mandatory disclosure The price of the asset with no, one or two ratings is given by (7), (8) or (10). All three prices are either constant or decreasing in the complexity of the asset. Hence, given the opportunity to design the asset, the issuer who is required to reveal every shadow rating he acquires prefers a less complex asset.

Voluntary disclosure The drawback that complex assets are riskier to investors and are, therefore, less profitable for asset issuers is still present with voluntary disclosure. But that drawback may now be offset by the following benefit. If ratings are drawn from a higher-variance distribution, their maximum will be higher, on average; that makes ratings shopping more profitable. Complex assets offer a broader menu of ratings for the issuer to pick from. When this advantage is large, the asset issuer will prefer to make his asset more complex.

Whether higher asset complexity is desirable depends on parameter values. Figure 3 illustrates a case where, for small levels of asset complexity, profits rise with added complexity. In this region with average ratings realizations, the issuer acquires only one shadow rating. But an issuer with a lower-than-average first rating acquires a second rating; if the two ratings are sufficiently far apart, he discloses the larger of the two. The fact that profit increases in σ_θ indicates a motive to design more-complex assets, which exacerbates the bias in disclosed ratings.

5 Evaluating Policy Recommendations

This section uses the model framework to explore the pros and cons of three proposed reforms: switching to investor-initiated ratings, increasing competition in the market for ratings, and reforming risk-management regulation.

5.1 Investor-initiated ratings

One possible solution to the problem of ratings bias is to replace issuer-initiated ratings with investor-initiated ratings. We show that even though some investors { those subject to the investment-grade securities regulations { would prefer biased ratings, they cannot shop for ratings. However, an investor-initiated market

for ratings may provide too little or even no information. To make this argument, we use the well-known model of information acquisition formulated by Grossman and Stiglitz (1980). We then add to that framework a requirement that some investors buy "investment grade" assets with sufficiently high ratings, as well as a market price, rather than a fixed cost, for information. Since the Grossman-Stiglitz framework is not our innovation, we describe its setup and main results briefly and refer the reader to Appendix A.2 and the original paper for more details.

5.1.1 A model of investor information acquisition

Investors choose whether or not to buy a single rating. In order to ensure that prices are not perfectly revealing, we will modify the assumption that the supply is known and fixed. We will instead assume that the issuer of the asset is endowed with $x + \tilde{x} = \int_i q_i^0 di$ shares of the asset. It is partly random: $x \sim N(0, \sigma_x)$. This randomness keeps investors from being able to free-ride on other investors' information.

We also change the specification of prior beliefs slightly to make the expressions simpler: The asset payoff[®] u is the sum of its rating θ , and the noise in the rating ϵ : $u = \theta + \epsilon$. Prior beliefs are that $\theta \sim N(u, \sigma_\theta^2)$ and $\epsilon \sim N(0, \sigma_\epsilon^2)$, where θ and ϵ are independent. Thus, we can write the prior belief about the payoff[®] as $u \sim (u, \sigma_\theta^2 + \sigma_\epsilon^2)$.

Let λ denote the fraction of investors that decide to buy a rating. A rating costs χ . Since investors are ex-ante identical, all investors buy information ($\lambda = 1$), all investors do not buy information ($\lambda = 0$) or all are indifferent. An indifference condition pins down equilibria where $\lambda \in (0, 1)$.

Model solution Since investors have the same utility function as in the previous section, their equilibrium asset demands satisfy the same first-order condition (equation 4). However, the information structure has changed. An informed investor, one who has observed the asset's rating θ , has posterior beliefs that the asset payoff[®] is distributed $u \sim N(\theta, \sigma_\epsilon^2)$. Substituting this posterior mean and variance into the first-order condition yields informed investors' asset demand:

$$q^I = \frac{1}{\rho} \sigma_\epsilon^{-2} (\theta - pr). \quad (25)$$

For uninformed investors, asset demand depends on information inferred from the asset price p . At the same time, the price depends on investors' demand. This is a fixed-point problem.

Grossman and Stiglitz (1980) show that the uninformed investors' optimal portfolio is

$$q^U = \frac{1}{\rho} \frac{\mu - pr}{\sigma_{\theta|p}^2 + \sigma_\epsilon^2}. \quad (26)$$

where

$$\sigma_{\theta|p}^2 \equiv V[\theta|\mathbf{t}, p] = \left[\sigma_{\theta}^{-2} + \left(\frac{\lambda \sigma_{\epsilon}^{-2}}{\rho} \right)^2 \sigma_x^{-2} \right]^{-1} \quad (27)$$

and the posterior mean $\hat{\mu}$ is given by equation (33) in appendix A.2.

Finally, the ratio of informed investors' expected utility to uninformed investors' expected utility, before accounting for information cost, is

$$\frac{E[U^I]}{E[U^U]} = \left(\frac{\sigma_{\epsilon}^2}{\sigma_{\epsilon}^2 + \sigma_{\theta|p}^2} \right)^{1/2}.$$

This ratio is less than one because utility is negative. Informed investors' utility is higher when it is less negative and, therefore, smaller in absolute value. In equilibrium, either there must be a corner solution $\lambda = \{0, 1\}$ or the expected net benefit of information ($E[U^I] - E[U^U]$) must equal the expected utility cost $-e^{\rho\chi} E[U^I]$. Thus, the condition for an interior equilibrium is $(\sigma_{\epsilon}^2 + \sigma_{\theta|p}^2)^{1/2} / \sigma_{\epsilon} = e^{\rho\chi}$.

5.1.2 Can investors acquire biased ratings?

Suppose that some subset of investors can buy only assets that are "investment grade." That means that the asset's rating surpasses a threshold ($\theta > \theta^{INV}$). This group of investors then achieves higher expected utility if they obtain an upward-biased rating: Let ϕ be the cumulative probability density function for the unbiased rating θ . Suppose that, instead of θ , the rating agency, with the knowledge of the investor, issues a rating $\theta + \epsilon$. Then, with probability $\phi(\theta^{INV} - \epsilon)$, the investor gets a rating that is too low; he cannot invest, and gets no income from the risky security. With probability $1 - \phi(\theta^{INV} - \epsilon)$, the investor can invest and chooses the optimal portfolio in equation (25). This informed investor's expected utility is proportional to

$$U^{bias} \propto - \int_{\theta^{INV} - \epsilon}^{\infty} E \left[\exp \left[-\frac{1}{2} (\theta - pr)^2 \sigma_{\epsilon}^{-2} \right] \right] d\theta. \quad (28)$$

Since the bias ϵ enters only in the bounds of integration and expands them, and since this is a negative number times an integrand that is strictly positive, ratings bias unambiguously increases expected utility.

Investor-initiated ratings bias has limited price impact If the investors shop for ratings, it is in order to find a rating agency that gives the asset an investment-grade rating. But once the investor finds that the asset is in his feasible investment set, he should use all available information to determine the optimal bidding strategy for the asset. Since investors use all available information in forming their asset demands, the only price effect of ratings bias is to raise the price of some assets that are not truly investment-grade to their price in the absence of regulation.

5.1.3 Downsides to investor-initiated ratings

Investor-initiated ratings can create two other problems due to information-market externalities: information leakage and market collapse due to demand complementarity. Since information requires a fixed cost to discover and is free (or, at least, quite cheap) to replicate, efficiency dictates that a discovered piece of information should be distributed to every asset investor so that all investors benefit from lower asset payoff risk. Yet, when investors have to pay for ratings themselves, either no investors or too few investors may end up being informed.

Information leakage One reason why asset investors may decide not to buy information is that they can partially free-ride on others' information. The price of the asset will depend on what informed investors know. While an uninformed investor cannot literally observe the price before bidding, he can condition on the price when submitting his menu of bids. When deciding on the quantity he associates with each realized price, the investor asks himself, "If that is the realized price, what would it tell me about what the informed investors have learned?" In this way, the uninformed investor can use information contained in prices to free-ride on what others have learned.

This information-leakage effect shows up in the model's solution. An increase in the number of informed investors λ reduces the posterior uncertainty of uninformed investors, conditional on the price level $\sigma_{\theta|p}$ (equation (27)). Therefore, it reduces the benefit of acquiring information (equation 28). As Grossman and Stiglitz (1980) point out, this is a form of strategic substitutability that makes it unlikely that investors will ever all choose to be informed.

Complementarity in information demand and market collapse Endogenizing the price of ratings so that what each investor pays $\chi(\lambda)$ depends on how many purchase the information introduces a complementarity that can collapse the market for information. Market collapse results when no investor buys a rating because no other investors are buying them. Such a collapse can arise in situations where the asset issuer would be willing to provide information to all investors.

Instead of assuming that ratings are provided at a fixed cost, we consider a profit-maximizing information-production sector. The sector has three crucial features: First, information can be produced with a fixed-cost technology. A rating θ can be discovered by any agent for a fixed cost c , which is the cost to the rating agency of collecting information to construct the rating (measured per capita). Once the rating is determined, it is costless to replicate it and sell it to multiple investors. Each investor pays the rating agency a price χ to observe its rating. Second, reselling purchased information is forbidden. The realistic counterpart to this assumption is intellectual property law that prohibits copying a publication and re-distributing it for profit. Third, there is free entry. Any agent can discover information at any time, even after other information producers have announced their prices χ . That information markets are competitive is crucial. The exact

market structure is not.¹⁷

Lemma 3 *The equilibrium price for information $\chi(\lambda)$ is decreasing in the quantity of information sold λ . Specifically, $\chi(\lambda) = c/\lambda$.*

With an exogenous cost c , there is a unique equilibrium. But with the endogenous price for information in an information market, there can be three equilibria. If $(\sigma_\epsilon^2 + \sigma_{\theta|p}^2)/\sigma_\epsilon^2)^{1/2} < e^{\rho c}$, then $\lambda = 0$ is an equilibrium. An investor who wants to acquire information must pay the entire fixed cost c for that information. This cost (in utility terms) exceeds the benefit of information. But it may also be the case that there are an additional two values of $\lambda \in (0, 1]$ such that $((\sigma_\epsilon^2 + \sigma_{\theta|p}^2)/\sigma_\epsilon^2)^{1/2} = e^{\rho c/\lambda}$. The lower of the two solutions will be an unstable equilibrium, but the higher one will be stable. As λ increases, the cost of information drops precipitously. At the same time, the information content of the asset's price $1/\sigma_{\theta|p}^2$ rises gradually, which reduces the benefit of information.

The problem arises when the equilibrium is the no-information ($\lambda = 0$) one. No investors buy information because no other investors are buying information, and if a given investor buys information by himself, he would have to bear the entire fixed cost c of discovering that information. In such a situation, the biased information provided by the asset-issuer-initiated ratings could result in more-accurate asset prices than could no information at all.

5.2 Increasing competition among rating agencies

Because market failures are often associated with a lack of free competition, policy makers have taken measures to increase competition. While this might cure some problems with ratings provision, it does not remedy ratings shopping. In fact, it worsens the problem.

When the issuer shops for ratings, the more draws the issuer can observe before choosing a rating $\{$ i.e., the larger the number of rating agencies $\}$ the higher this bias will be.

Proposition 4 *If the asset issuer will disclose only the most favorable rating, then increasing the number of ratings agencies will (weakly) increase the bias of the disclosed rating.*

This result follows from the simple observation that the more rating agencies available, the greater the possibilities of ratings shopping. Of course, having more rating agencies does not ensure that an asset issuer will observe more shadow ratings. If not, then the bias will stay constant. However, if some issuers prefer to obtain more shadow ratings than what were previously available to them, increasing the number of agencies will increase the number of observed ratings and the bias from shopping for the best one. It is also possible that the price of shadow ratings falls due to higher competition, encouraging asset issuers to sample more ratings, which would increase ratings bias even more.

¹⁷Veldkamp (2006) analyzes Cournot and monopolistic competition markets for information. All three markets produce information prices that decrease in demand.

5.3 The effect of risk-management legislation

Another target for criticism in the ratings scandal has been the role of risk-management rules. Many banks and pension funds are required to hold only investment-grade securities. These are assets that earn sufficiently high ratings from one of the nationally-recognized statistical ratings organizations (NRSRO's) (this group includes Moody's, S&P and Fitch). This rule puts an enormous amount of pressure on asset issuers to ensure that their assets achieve this rating. Without it, the pool of potential investors is considerably smaller, and the asset's prices will be considerably lower.

Repealing the investment-grade securities regulations, alone, will not solve the problem of ratings shopping. With investor-initiated ratings, the bias arose without the regulation. However, it is likely that this regulation further encouraged ratings shopping by increasing the payoff from acquiring a high rating, for a given level of ratings uncertainty.

5.4 Mandatory disclosure laws

Perhaps an obvious suggestion is to mandate disclosure of all ratings. While, in theory, that is a cure, in practice, it is difficult to directly regulate the transmission of information. For example, the line between informal advice and a rating can be easily blurred. Prohibiting a discussion of how various assets might be rated if they were issued could easily be ruled an infringement on free speech. An additional problem is that when undesirable ratings are proposed, the asset in question is frequently restructured. A tiny change in asset structure would make the previous rating no longer applicable and could effectively hide that rating.

6 Conclusions

Examining common arguments about how ratings may have distorted credit-derivative prices exposed logical gaps. But it also suggested a coherent story about why ratings bias might have emerged and why investors could not use past data to detect it. Developing a model of the market for ratings in which asset issuers can shop for ratings revealed that, under some circumstances, an increase in asset complexity could generate ratings bias. Solving the model also delivered additional insights. It revealed a feedback effect whereby an increase in asset complexity prompted ratings shopping, which gave issuers an incentive to structure assets that are even more complex. It also illustrated how more competition in ratings markets could make the distortions in ratings even more severe. An extension of the model in which investors purchase ratings, uncovered multiple equilibria. This taught us that a move to investor-initiated ratings could collapse the market for ratings, leaving investors with information that is even less reliable.

None of the policy options we examined were without their drawbacks. Yet, the model points to two possible solutions to the ratings-bias problem. Investor-initiated ratings are a cure for bias. The problem

of market collapse is mitigated when the investors are large players in their markets who find it valuable to purchase information, even if other investors do not. Since most complex credit products are purchased by institutional investors, rather than by households, the large-investor assumption might not be a bad one. The second possible solution is to have one rating agency, a regulated monopoly, that rates every bond. If every asset has one, and only one, rating, shopping is not possible. Of course, regulated monopolies have little incentive to provide reliable information. The choice between these two options is a quantity-versus-quality choice. An investor-initiated system will produce less information, but it will be more reliable. Finally, it could be that policy reform is not needed. A key assumption of the model is that investors are unaware of the increase in asset complexity and the resulting ratings bias. But the crisis has drawn attention to ratings practices. If investors who observe a small number of ratings infer that some ratings are being withheld, then ratings shopping will no longer be profitable, and additional regulation to cure the problem would be unnecessary.

Many markets supply information or certification services: Academic testing services, appraisals, or job head-hunters are just a few examples. (See Bar-Isaac, Caruana and Cunat (2008).) Our paper raises the question: What determines the quality of the information produced? We point out that not only does the nature of the good being sold affect the information available about it, but also that the nature of the evaluated products may change to game the ratings system, possibly to disastrous effect.

References

- Banerjee, Snehal**, \Learning from Prices and Dispersion in Beliefs," 2007. Northwestern University working paper.
- Bar-Isaac, Heski, Guillermo Caruana, and Vicente Cunat**, \Information Gathering Externalities in Product Markets," 2008. Working Paper.
- Barlevy, Gadi and Pietro Veronesi**, \Information Acquisition in Financial Markets," *Review of Economic Studies*, 2000, 67, 79{90.
- Becker, Bo and Todd Milbourn**, \Reputation and Competition: Evidence from the Credit Rating Industry," 2008. HBS Finance working paper 09-051.
- Benabou, Roland and Guy Laroque**, \Using Privileged Information to Manipulate Markets: Insiders, Gurus, and Credibility," *Quarterly Journal of Economics*,, 1992.
- Bolton, Patrick, Xavier Freixas, and Joel Shapiro**, \Conflicts of interest, information provision, and competition in the financial services industry," *Journal of Financial Economics*, 2007.
- , —, and —, \The Credit Ratings Game," 2008. Working paper.
- Brunnermeier, Markus**, *Asset Pricing under Asymmetric Information: Bubbles, Crashes, Technical Analysis and Herding*, first ed., Oxford University Press, 2001.
- Bullard, James, George Evans, and Seppo Honkapohja**, \Near-Rational Exuberance," 2005. Working Paper.
- Damiano, E, H Li, and W Suen**, \Credible Ratings," *Theoretical Economics*, 2008, 3, 325{365.
- DeMarzo, Peter, Dimitri Vayanos, and Jeff Zwiebel**, \Persuasion Bias, Social Influence, and Unidimensional Opinions," *Quarterly Journal of Economics*, 2003, 118, 909{968.
- Esponda, Ignacio**, \Behavioral Equilibrium in Economies with Adverse Selection," *American Economic Review*, 2008, 98(4), 1269{91.
- Eyster, E and M Rabin**, \Cursed Equilibrium," *Econometrica*, 2005, 73(5).
- Farhi, Emmanuel, Josh Lerner, and Jean Tirole**, \Fear of Rejection? Tiered Certification and Transparency," 2008. Harvard University working paper.
- Faure-Grimaud, A, E Peyrache, and L Quesada**, \The Ownership of Ratings," 2007. FMG Discussion Papers dp590, Financial Markets Group.
- Grossman, Sanford**, \The Informational Role of Warranties and Private Disclosure About Product Quality," *Journal of Law and Economics*, 1981, 24, 461{489.
- and **Joseph Stiglitz**, \On the Impossibility of Informationally Efficient Markets," *American Economic Review*, 1980, 70(3), 393{408.
- Hellwig, Martin**, \On the Aggregation of Information in Competitive Markets," *Journal of Economic Theory*, 1980, 22, 477{498.
- Hu, Jian**, \Assessing the Credit Risk of CDOs Backed by Structured Finance Securities: Rating Analysts' Challenges and Solutions," 2007. Working Paper, Moody's Investors' Service.
- Jovanovic, Boyan**, \Truthful Disclosure of Information," *Bell Journal of Economics*, 1982, 13, 36{44.

- Kotz, S., N. Balakrishnan, and N.L. Johnson**, *Continuous Multivariate Distributions, Volume 1, Models and Applications*, second ed., Oxford University Press, 2000.
- Mason, Joseph R. and Josh Rosner**, "Where Did the Risk Go? How Misapplied Bond Ratings Cause Mortgage Backed Securities and Collateralized Debt Obligation Market Disruptions," 2007. SSRN Working Paper #1027475.
- Milgrom, Paul**, "Good news and Bad news: Representation Theorems and Applications," *The Bell Journal of Economics*, 1981, 12, 380{391.
- Morgan, J and P Stocken**, "An analysis of stock recommendations," *The Rand Journal of Economics*, 2003.
- Peress, Joel**, "Wealth, Information Acquisition and Portfolio Choice," *The Review of Financial Studies*, 2004, 17(3), 879{914.
- Reny, P and M Perry**, "Toward a Strategic Foundation of Rational Expectations Equilibrium," *Econometrica*, 2006, 74 (5), 1231{1269.
- Sangiorgi, Francesco, Jonathan Sokobin, and Chester Spatt**, "Credit-Rating Shopping, Selection and the Equilibrium Structure of Ratings," 2008. Carnegie Mellon working paper.
- Shavell, S**, "Acquisition and Disclosure of Information Prior to Sale," *Rand Journal of Economics*, 1994, 25, 20{36.
- Veldkamp, Laura**, "Media Frenzies in Markets for Financial Information," *American Economic Review*, 2006, 96(3), 577{601.
- White, Lawrence**, "A New Law for the Bond Rating Industry," *Regulation*, 2007, 30(1), 48{52.

A Appendix: Proofs and Derivations

A.1 Computation Details for the Disclosure Decision 2 versus 0

The question is the following: Is it possible to have $\Pi_{D=2}(\bar{\theta}, \underline{\theta}) - \Pi_{D=0} > 0$ when $\bar{\theta} \leq a \equiv \frac{r(1+\sigma_u^{-2}\sigma_\theta^2)}{x}\chi + \bar{u} - \rho x \sigma_u^2$? When the maximum of the two draws, $\bar{\theta}$, is below $a = \frac{r(1+\sigma_u^{-2}\sigma_\theta^2)}{x}\chi + \bar{u} - \rho x \sigma_u^2$ then the second highest draw, $\underline{\theta}$, will also be necessarily below $\frac{r(1+\sigma_u^{-2}\sigma_\theta^2)}{x}\chi + \bar{u} - \rho x \sigma_u^2$. Then, when $\bar{\theta} \leq a$, the highest possible sum for $\bar{\theta} + \underline{\theta}$ is $2a$. We now show that $\Pi_{D=2}(\bar{\theta}, \underline{\theta}) - \Pi_{D=0} > 0$ cannot be satisfied even when $\bar{\theta} + \underline{\theta}$ is $2a$.

To see this, note that

$$\begin{aligned}
& \left[\frac{1}{r} \frac{\sigma_u^{-2}\bar{u} + \sigma_\theta^{-2}2 \cdot \left(\frac{r(1+\sigma_u^{-2}\sigma_\theta^2)}{x}\chi + \bar{u} - \rho x \sigma_u^2 \right) - \rho x}{\sigma_u^{-2} + 2\sigma_\theta^{-2}} - \frac{1}{r}(\bar{u} - \rho \sigma_u^2 x) \right] x - 2\chi \\
&= \left[\frac{1}{r} \frac{\sigma_u^{-2}\bar{u} + \left(\frac{\sigma_\theta^{-2}2r(1+\sigma_u^{-2}\sigma_\theta^2)}{x}\chi + \sigma_\theta^{-2}2\bar{u} - \sigma_\theta^{-2}2\rho x \sigma_u^2 \right) - \rho x}{\sigma_u^{-2} + 2\sigma_\theta^{-2}} - \frac{1}{r}(\bar{u} - \rho \sigma_u^2 x) \right] x - 2\chi \\
&= \left[\frac{1}{r} \frac{\sigma_u^{-2}\bar{u} + \left(\frac{\sigma_\theta^{-2}2r(1+\sigma_u^{-2}\sigma_\theta^2)}{x}\chi + \sigma_\theta^{-2}2\bar{u} - \sigma_\theta^{-2}2\rho x \sigma_u^2 \right) - \rho x}{\sigma_u^{-2} + 2\sigma_\theta^{-2}} - \frac{1}{r}(\bar{u} - \rho \sigma_u^2 x) \right] x - 2\chi \\
&= \left[\frac{1}{r} \frac{\sigma_u^{-2}\bar{u} + \sigma_\theta^{-2}2\bar{u} - \sigma_\theta^{-2}2\rho x \sigma_u^2 - \rho x}{\sigma_u^{-2} + 2\sigma_\theta^{-2}} + \frac{\sigma_\theta^{-2}2r(1+\sigma_u^{-2}\sigma_\theta^2)}{x(\sigma_u^{-2} + 2\sigma_\theta^{-2})}\chi - \frac{1}{r}(\bar{u} - \rho \sigma_u^2 x) \right] x - 2\chi \\
&= \left[\frac{1}{r} \frac{\sigma_u^{-2}\bar{u} + \sigma_\theta^{-2}2\bar{u} - \sigma_\theta^{-2}2\rho x \sigma_u^2 - \rho x}{\sigma_u^{-2} + 2\sigma_\theta^{-2}} - \frac{1}{r}(\bar{u} - \rho \sigma_u^2 x) \right] x + \left(\frac{2(\sigma_\theta^{-2} + \sigma_u^{-2})}{(\sigma_u^{-2} + 2\sigma_\theta^{-2})} - 2 \right) \chi \\
&= \left[\frac{1}{r} \left[\frac{\sigma_u^{-2}\bar{u} + \sigma_\theta^{-2}2\bar{u}}{\sigma_u^{-2} + 2\sigma_\theta^{-2}} + \frac{1 - \rho x \sigma_u^2 (\sigma_\theta^{-2}2 + \sigma_u^{-2})}{\sigma_u^{-2} + 2\sigma_\theta^{-2}} \right] - \frac{1}{r}(\bar{u} - \rho \sigma_u^2 x) \right] x + \left(\frac{2(\sigma_\theta^{-2} + \sigma_u^{-2})}{(\sigma_u^{-2} + 2\sigma_\theta^{-2})} - 2 \right) \chi \\
&= \left[\frac{1}{r} [\bar{u} - \rho \sigma_u^2 x] - \frac{1}{r}(\bar{u} - \rho \sigma_u^2 x) \right] x + \left(\frac{2(\sigma_\theta^{-2} + \sigma_u^{-2})}{(\sigma_u^{-2} + 2\sigma_\theta^{-2})} - 2 \right) \chi \\
&= \left(\frac{2(\sigma_\theta^{-2} + \sigma_u^{-2})}{(\sigma_u^{-2} + 2\sigma_\theta^{-2})} - 2 \right) \chi < 0.
\end{aligned}$$

A.2 Solution to the investor-initiated ratings model

The price of the risky asset is determined by the market-clearing condition

$$\lambda q^I + (1 - \lambda)q^U = x + \bar{x}. \quad (29)$$

Uninformed investors combine their prior belief that $\theta \sim N(\bar{u}, \sigma_\theta^2)$ and their signal from the price $(p - A)/B \sim N(\theta, (C/B)^2 \sigma_x^2)$ to form their posterior belief: $\theta \sim N(\hat{\mu}, \sigma_{\theta|p}^2)$ where the posterior variance is given by equation (27).

Grossman and Stiglitz (1980) conjecture that prices take the linear form $p = A + B\theta + Cx$, determine what this price implies for risky-asset demands, substitute those demand functions into the market-clearing conditions, and match coefficients to verify the hypothesis. Substituting asset demands (25) and (26) in to the market-clearing condition (29), yields

$$\lambda \sigma_\epsilon^{-2}(\theta - pr) + (1 - \lambda) \frac{\hat{\mu} - pr}{\sigma_{\theta|p}^2 + \sigma_\epsilon^2} = \rho(\bar{x} + x).$$

Substituting in for $\sigma_{\theta|p}^2$ from (27) and $\hat{\mu}$ from (33) and collecting price terms reveals that the equilibrium price formula is linear in x and θ :

$$A = -\psi \left[\rho \bar{x} + (1 - \lambda) \frac{\hat{\sigma}_{\theta|p}^2}{\hat{\sigma}_{\theta|p}^2 + \sigma_\epsilon^2} \left(\frac{AB\sigma_x^2}{C^2} - \sigma_\theta^{-2}\mu \right) \right] \quad (30)$$

$$B = \psi \lambda \sigma_\epsilon^{-2} \quad (31)$$

$$C = -\psi \rho \quad (32)$$

where $\psi = [\lambda\sigma_\epsilon^2 r + (1 - \lambda)\frac{\hat{\sigma}_{\theta|p}^2}{\sigma_{\theta|p}^2 + \sigma_\epsilon^2} \frac{\lambda\sigma_\epsilon^{-2}\sigma_x^2}{\rho C}]^{-1}$. Using the fact that $(p - A)/B$ is an unbiased signal about θ , with variance $(C/B)^2\sigma_x^2$, we can then use Bayes' law to compute the posterior expectation of uninformed investors:

$$\hat{\mu} = \sigma_{\theta|p}^2 \left[\sigma_{\theta}^{-2} \hat{\mu} + \left(\frac{B}{C} \right)^2 \sigma_x^{-2} \left(\frac{p - A}{B} \right) \right]. \quad (33)$$

Expected Utility for Uninformed Investors To compute the value of information, we proceed in three steps, using the law of iterated expectations: $E[e^{-\rho W}] = E[E[e^{-\rho W}|E[u], p] | p]$. First, take the first expectation, over u . Using the formula for the mean of a log-normal, we get

$$E[u^U | pr] = -\exp\left[-\frac{1}{2} \frac{(\hat{\mu} - pr)^2}{\sigma_\epsilon^2 + \sigma_{\theta|p}^2}\right].$$

There is a second expectation over pr that we haven't taken. It turns out that we won't need to. So, we leave this line of argument for now.

Expected Utility for Informed Investors For informed investors, following the same steps yields:

$$= E[e^{-\frac{1}{2}(\theta - pr)' \sigma_\epsilon^{-2}(\theta - pr)}],$$

where θ replaced $\hat{\mu}$ as the conditional mean and σ_ϵ^2 is now the conditional variance. The next step is to take an expectation over θ , but not pr .

Applying the formula for a moment-generating-function of a quadratic normal:

$$E[U^I | rp] = -|I - 2\sigma_{\theta|p}^2(-\frac{1}{2})\sigma_\epsilon^{-2}|^{-1/2} \exp\left[\frac{1}{2}(\hat{\mu} - pr)^2 \sigma_\epsilon^{-4}(I + \sigma_{\theta|p}^2 \sigma_\epsilon^{-2})^{-1} \sigma_{\theta|p}^2 - \frac{1}{2}(\hat{\mu} - pr)^2 \sigma_\epsilon^{-2}\right].$$

Algebraic manipulation yields

$$E[U^I | rp] = \left(\frac{\sigma_\epsilon^2}{\sigma_\epsilon^2 + \sigma_{\theta|p}^2}\right)^{1/2} \exp\left[-\frac{1}{2} \frac{(\hat{\mu} - pr)^2}{\sigma_\epsilon^2 + \sigma_{\theta|p}^2}\right] \quad (34)$$

$$= \left(\frac{\sigma_\epsilon^2}{\sigma_\epsilon^2 + \sigma_{\theta|p}^2}\right)^{1/2} E[U^U | rp]. \quad (35)$$

This ratio of variances is a known quantity when information is acquired. All agents can infer the equilibrium strategies of other agents and deduce how much information will be revealed through the price level. Since the two expected utilities are related by a known constant, their unconditional expectations $E[U^I]$ and $E[U^U]$ must be proportional, as well.

A.3 Proof of Lemma 3

Let $d_{it} = 1$ if agent i decides to discover information in period t , and $d_{it} = 0$ otherwise. Let per-capita demand for information with price χ_{it} , given all other posted prices χ_{-it} , be $I(\cdot, \cdot)$. Then, the objective of the information producer is to maximize profit:

$$\max_{d_{it}, \chi_{it}} d_{it}(\chi_{it} I(\chi_{it}, \chi_{-it}) - c). \quad (36)$$

Suppose that the equilibrium information price was above average cost $\chi > c/\lambda$. Then, an alternate supplier could enter the market with a slightly lower price, and make a profit. If a supplier set price below marginal cost, it would make a loss. This strategy would be dominated by no information provision. If there are two or more suppliers, then either price is above marginal cost, which can't be an equilibrium by the first argument, or both firms price at (or below) marginal cost, split the market, and make a loss, which is dominated by exit.

B Appendix: Rating Agencies' Incentive to Bias Ratings

Here, we sketch a model in which a rating agency can choose to bias ratings if it finds it profitable to do so. The agency's utility depends on its profits and a reputation cost that is a quadratic function of the distance between its forecast and the true asset payoff.

Suppose that once a rating is ordered, it is paid for. Then, the rating agency's payoff is given by:

$$U^r = \chi + \tilde{\chi} - c - \alpha(\theta_i - u)^2,$$

where c stands for the cost of producing a rating. Suppose that the investor of a rating can observe the rating before purchasing it. Then, in this case, the profit of the rating agency depends on the probability that the asset issuer purchases its rating $\pi(\theta_i)$ and the ratings price χ

$$U^r = \pi(\tilde{\theta}_i)(\chi + \tilde{\chi}) - \alpha(\tilde{\theta}_i - u)^2 - c. \quad (37)$$

the distribution of the highest-order statistic, $\theta_i = r(\theta_i)$. When the issuer buys the highest of all ratings, then

$$\pi_i(\tilde{\theta}_i) = \text{prob}(\tilde{\theta}_i = \max_{j \in I} \tilde{\theta}_j).$$

There are two equivalent ways to model ratings inflation: One way is to assume that rating agencies draw a rating θ_i from an unbiased distribution, but they report a different rating $\tilde{\theta}_i = r(\theta_i)$. The other is to model rating agencies that draw from biased distributions. We sketch the first approach.

Assuming that all rating agencies use the same technology to produce ratings (draw ratings from the same distribution), and that all I rating agencies use the same monotonic (strictly increasing in θ_i) reporting strategy, then

$$\pi(\theta_i) = F^{I-1}(r^{-1}(\tilde{\theta}_i)),$$

and the payoff of the rating agency is given by

$$U^r = F^{I-1}(r^{-1}(\tilde{\theta}_i))(\chi + \chi) - \alpha(\theta_i - u)^2 - c.$$

Heuristic derivation of the equilibrium reporting strategy: