The Journal of FINANCE
The Journal of THE AMERICAN FINANCE ASSOCIATION

THE JOURNAL OF FINANCE • VOL. LXVII, NO. 1 • FEBRUARY 2012

The Credit Ratings Game

PATRICK BOLTON, XAVIER FREIXAS and JOEL SHAPIRO*

ABSTRACT

The collapse of AAA-rated structured finance products in 2007 to 2008 has brought renewed attention to conflicts of interest in credit rating agencies (CRAs). We model competition among CRAs with three sources of conflicts: (1) CRAs conflict of understating risk to attract business, (2) issuers' ability to purchase only the most favorable ratings, and (3) the trusting nature of some investor clienteles. These conflicts create two distortions. First, competition can reduce efficiency, as it facilitates ratings shopping. Second, ratings are more likely to be inflated during booms and when investors are more trusting. We also discuss efficiency-enhancing regulatory interventions.

[The investment] could be structured by cows and we would rate it
—Analyst at one of the main credit rating agencies in an e-mail referring to structured finance products, April 5, 2007.[1]

THE ANALYST IN THE above statement refers to a key dilemma for credit rating agencies (CRAs): How should they act when their principal source of revenue comes from the firms whose products they are rating? This potential source of conflict has repeatedly been brought to the public's (and regulators') attention, particularly following the East Asian Financial Crisis (1997) and in the aftermath of the failures of Enron (2001) and Worldcom (2002), but it has never been so salient as during the recent financial crisis. Indeed, while CRA profits exploded with the growth of structured finance products (Moody's profits, e.g., tripled between 2002 and 2006[2]), the large number of downgrades of these

*Bolton is with Columbia University; Freixas is with Universitat Pompeu Fabra; and Shapiro is with University of Oxford. Bolton, Freixas, and Shapiro are also affiliated with CEPR. We thank Adam Ashcraft, Heski Bar-Isaac, Marco Becht, Bo Becker, Itay Goldstein, Doh-Shin Jeon, Jens Josephson, Navin Kartik, Marco Pagano, Ailsa Röell, Francesco Sangiorgi, David Skeie, Chester Spatt, James Vickery, Mungo Wilson, and audiences at AFA 2010, Bank of England, Bergen, BIS, Bocconi, CEPR conference on Transparency, Disclosure and Market Discipline in Banking Regulation, FIRS 2009, Hunter, NBER Workshop on the Economics of Credit Rating Agencies, NBER Securitization Working Group, NY Fed, NYU-Stern, SEC, Tilburg, and UPF for helpful discussions. We also thank the Editor (Cam Harvey), an Associate Editor, and the referee for detailed comments and suggestions. Shapiro gratefully acknowledges financial support from the Spanish government under SEJ2006-09993 and from the Fundacion Ramon Areces. Freixas acknowledges financial support from The Barcelona GSE Research Network and the Generalitat de Catalunya.

[1] Securities and Exchange Commission (2008, p. 12).

[2] "Triple-A-Failure," by Roger Lowenstein, *New York Times Magazine*, April 27, 2008. Moody's profits are the easiest of the CRAs to measure since they are a public stand-alone company. "Moody's

securities from 2007 onwards has fostered suspicion that ratings standards had been relaxed during the boom years. Further along with these allegations of possible conflicts of interest for CRAs, many commentators have also reproved (institutional) bond investors for their excessive reliance on ratings, and for not doing their homework in independently assessing default risk. The combination of CRA reliance on fees from issuers, investors who were too trusting, and issuers looking to benefit from the mispricing of their issues could have led to substantial ratings inflation with important systemic consequences.

In this paper, we combine these elements in a model of credit ratings and CRA competition to analyze the equilibrium outcome of ratings and the efficiency consequences of possible equilibrium ratings inflation. The model gives rise to two fundamental equilibrium distortions. First, competition among CRAs may reduce market efficiency since it facilitates ratings shopping by issuers and results in excessively high *reported* ratings. We show in particular that, as a result of issuer shopping, efficiency may be higher under a monopoly CRA than under a duopoly despite the potential for the increased informativeness of two ratings. Second, CRAs are more prone to inflate ratings during booms, when there is a larger clientele of investors in the market who take ratings at face value and when the risks of failure that could damage CRA reputation are lower.

The key building blocks of our model are as follows:

1. *Issuer payments for ratings*: In practice, CRA fees involve both a fee at the time of issuance and an annual fee for as long as the issue is outstanding. Importantly, while CRAs have list price schedules, they may renegotiate fees with regular customers (White (2002)).[3] In addition, CRAs offer related consulting services, such as prerating assessments.
2. *Issuer shopping for ratings*: In practice, as in our model, an issuer pays a CRA only if it asks the CRA to make the rating public.[4] Also, if an issuer is unhappy with a rating, it may solicit another one.[5]
3. *CRA credit models may vary in precision*: We consider CRA credit risk models that provide imperfect assessments of default risk. As Deven Sharma,

operating margins exceeded 50% for the past 6 years, three to four times those of Exxon Mobil Corp., the world's biggest oil company." "Bringing Down Wall Street as Ratings Let Loose Subprime Scourge," by Elliot Blair Smith, www.bloomberg.com, September 24, 2008.

[3] The Securities and Exchange Commission found that, in a sample of subprime residential mortgage-backed security (RMBS) deals, 12 arrangers represented 80% of the business in both number and dollar volume, while for collateralized debt obligations (CDOs) of subprime deals, 11 arrangers accounted for 92% of the deals and 80% of the dollar volume (Securities and Exchange Commission (2008, p. 32)).

[4] "Typically the rating agency is paid only if the credit rating is issued" (Securities and Exchange Commission (2008, p. 9)).

[5] "Brian Clarkson, then president and chief operating officer of Moody's Investor's Service acknowledged that, 'There is a lot of rating shopping that goes on . . . What the market doesn't know is who's seen certain transactions but wasn't hired to rate those deals.'" "Bond-Rating Shifts Loom in Settlement; N.Y.'s Cuomo Plans Overhaul of How Firms Get Paid," by Aaron Lucchetti, *Wall Street Journal*, June 4, 2008.

President of Standard & Poor's (S&P), notes: "Events have demonstrated that the historical data we used and the assumptions we made significantly underestimated the severity of what has actually occurred." [6]

4. *CRAs can make "adjustments" to their credit risk model outputs*: As Griffin and Tang (2010) show in their study of structured product credit ratings, CRAs use *noisy* credit risk models, to which they make frequent adjustments before determining the final rating. Importantly for our analysis, the authors show that these adjustments tend to shift the rating upwards relative to the model-predicted rating.
5. *Reputation concerns for CRAs*: As rating agencies executives often argue, CRAs are concerned about maintaining their reputation for providing timely and accurate assessments of (changes in) default risk. Accordingly, we introduce in our model a reputation cost that CRAs incur in the event that an issue they rated highly ends up in default. Short-term gains from inflating an issue's quality can thus be smaller in our model than long-term reputation losses from jaded investors.
6. *Barriers to entry in the credit rating industry*: We confine our analysis to competition between two CRAs. However, it is possible, although somewhat tedious, to extend our analysis to the case of three CRAs, which broadly speaking is the current market structure in the credit rating industry. This high concentration of CRAs is a reflection of large barriers to entry into this industry. One "artificial" barrier was established by the Securities and Exchange Commission (SEC), in 1975 when it created the *Nationally Recognized Statistical Rating Organization* (NRSRO) designator, which indicates those CRAs whose ratings the SEC recognizes as being valuable for investment decisions. Although seven firms initially received this designation, mergers brought this number down to three (Standard & Poor's, Moody's, and Fitch) and the SEC had not admitted new firms until recently.[7] Since Congress, local governments, and regulatory agencies adopt this designation, White (2002, p. 52) argues that it has resulted in an "absolute barrier to entry." The extremely high profit margins of CRAs are also emblematic of a highly concentrated industry.
7. *Sophisticated and "trusting" investor clienteles*: Some of the potential investors in rated issues are sophisticated and understand a CRA's potential conflicts of interest; they are thus able to see through ratings inflation. However, a significant fraction (which may vary) of investors are trusting, in that they take the CRAs' ratings at face value. This coexistence of trusting and sophisticated investors may be due to different types of incentives to perform due diligence. Trusting investors, for example, may be pension fund managers, whose compensation only marginally depends on the ex post return of the assets they manage. Moreover, the more

[6] Testimony before the Committee on Oversight and Government Reform, United States House of Representatives, Deven Sharma, October 22, 2008.

[7] Until 2003, when the SEC gave Dominion the NRSRO designation. In 2005, A.M. Best received the designation, and in 2006 three more designations were given out (White (2010)).

complex the investments, the more costly it may be to uncover their value. Sophisticated investors, on the other hand, could be hedge funds, whose returns depend more directly on the profitability of the investment. Regulation that forces managers to only purchase investments with good ratings could also provide incentives to be trusting.[8] In a study of the CRA credit watch mechanism, Boot, Milbourn, and Schmeits (2006) model investors who take ratings at face value, calling them institutional investors. Similarly, Hirshleifer and Teoh (2003, p. 338) model investors with "limited attention and processing power." More generally, this allows for a rich and subtle interaction between two different investor clienteles (which seems of the essence for CRAs) and contributes to the literature on differences of opinion.[9]

Incorporating these key features into our model, we demonstrate under what situations ratings inflation is more likely to occur, what its impact on market efficiency is likely to be, and what the impact of regulatory proposals is likely to be. Furthermore, we examine empirical implications of the model and evidence from current studies on CRAs and structured finance products. We now summarize the main results.

Our most important result is that a duopoly ratings industry is generally less efficient than a monopoly. The reason is that, although in a duopoly investors could obtain more information, the issuer has more opportunities to shop for a good rating and to take advantage of trusting investors by only purchasing the best ratings. By extending the model to two periods to allow for endogenous reputation, we further show that the greater efficiency of a monopoly CRA holds for any parameter constellation. This result is consistent with the findings of Becker and Milbourn (2011), who show that the greater competitive threat posed by Fitch in the corporate bond market coincides with a deterioration in ratings quality.

We next show that CRAs may inflate the quality of the issuer's investment when there are more trusting investors in the market and/or when CRA expected reputation costs are lower. As these features are common to entire classes of issues of similar characteristics, ratings inflation is not just about idiosyncratic attributes of a single issuer but rather has systemic effects. In particular, during boom times, when more investors are trusting and the probability of getting caught is smaller, more ratings inflation is likely to occur. This result is consistent with the findings of Ashcraft, Goldsmith-Pinkham, and

[8] One might argue that investors were all sophisticated because they both originated and held these securities. However, this is not true empirically. From Lehman Brothers calculations in April 2008, we know that U.S. commercial and savings banks represented only 23% of the holders of nonagency AAA securities. ("Residential Credit Losses—Going into Extra Innings?" Lehman Brothers Securitized Products Research, 2008.) If we were generous, we might add broker dealers (who held 6.1%) to the list of possible originators. The main other holders were government-sponsored enterprises (GSEs) and the Federal Home Loan Bank (FHLB) System (18.8%), money managers (13.8%), insurance companies (7.6%), and overseas investors (25.2%).

[9] We provide a somewhat different (more institution-based) explanation for why differences of opinion arise (see Harrison and Kreps (1978) and Scheinkman and Xiong (2003)).

Vickrey (2009), who show that ratings of mortgage-backed securities (MBSs) were least accurate at the peak of the real estate boom. We also show that more precise CRA credit risk models enhance CRA payoffs from inflating their ratings as well as increase their probability of getting caught ex post, so that their overall effect is ambiguous.

We further show that, when an issuer is more important to a CRA, either because it is a repeat issuer or because it has larger issues, the CRA is more prone to inflate that issuer's ratings. This result is in line with the findings of He, Qian, and Strahan (2010), who show that CRAs rated large structured product issuers more favorably, and Faltin-Traeger (2009), who finds that repeat issuers are more likely to stick with the same CRA if they received a more favorable early rating.

Finally, we analyze reforms to the industry in the context of our model. The Cuomo plan, which is an agreement between New York State Attorney General Andrew Cuomo and the three main CRAs, requires that issuers pay CRAs for their rating up-front, not contingent on the report. In our model, this plan eliminates the incentives for CRAs to inflate ratings, but does not eliminate shopping. Mandating automatic disclosure of any ratings solicited is therefore necessary to get rid of the shopping distortion.[10] In addition, the up-front fees may undermine CRA incentives to invest in model accuracy and due diligence, making oversight on methodology potentially important.

Next, we offer a summary of the related theoretical literature (Section VI is dedicated to a discussion of the empirical evidence).

Related Theoretical Literature

There is a substantial literature on information intermediaries in both microeconomics and finance. The paper closest to ours is Mathis, McAndrews, and Rochet (2009), who examine the incentives of a monopoly CRA to inflate ratings in a model of endogenous reputation.[11] They find that reputation cycles may exist where a CRA builds up its reputation by relaying information accurately only to exploit this reputation later by collecting fees for inflated ratings. They also demonstrate that truth telling incentives are weaker when the CRA has more business from rating complex products. While their model endogenizes reputation, it restricts them to analyzing only a monopolist and to define a complex product simply as one in which the CRA's reputation is at stake. By making the large assumption that reputation is exogenous, we are able to examine the effects of competition and include a wealth of parameters on which we can perform comparative statics. Nevertheless, we endogenize reputation

[10] This regulation can only address explicit shopping. The implicit shopping problem would remain (see Sangiorgi, Sokobin, and Spatt (2009)).

[11] Strausz (2005) and Bar-Isaac and Shapiro (2011) also model endogenous reputation for information intermediaries. Strausz (2005) provides interesting insights in line with our findings, as he argues that honest certification has some of the characteristics of a natural monopoly. Bar-Isaac and Shapiro (2011) incorporate economic shocks and show that CRA accuracy may be countercyclical, which is also consistent with our results.

in a simple repeated game in Section VI to show that our results are indeed robust.

In the microeconomics literature, information intermediaries are modeled in acquiring and certifying information by committing to disclosure rules, as, for example, in Biglaiser (1993) and Lizzeri (1999). In contrast, CRAs do not commit to information disclosure rules and their incentives come from the possible reputation costs they incur when they provide inaccurate information. This is akin to the issues financial analysts face when they recommend stocks, as analyzed by Benabou and Laroque (1992) and Morgan and Stocken (2003). The model of Morgan and Stocken (2003) also addresses the issue of unverifiable information provision, when the certifier can lie but thereby incurs a lying cost (this problem is examined further in Kartik (2009), Kartik, Ottaviani, and Squintani (2007), and Ottaviani and Sorensen (2006)).

Although our signaling game is simpler in some respects, we extend the literature by examining how strategic contracting between the informed party (the CRA) and an interested party (the issuer in our case) can affect information revelation. Our problem is also related to the economics literature on strategic contracting when the information revealed affects a third party, which covers a wide number of microeconomics issues (see Inderst and Ottaviani (2009), Durbin and Iyer (2009), and Mariano (2008)). In Pagano and Volpin (2010), CRAs have no conflicts of interest, but can choose to be more or less opaque depending on what the issuer asks for. The authors show that, because of the existence of a winner's curse, opacity can enhance liquidity in the primary market but may cause a market freeze in the secondary market.

In Bolton, Freixas, and Shapiro (2007), we analyze a situation of strategic contracting where the informed parties (banks) set prices for their products at the same time as they provide recommendations about them to uninformed investors. We show that competition unambiguously reduces banks' incentives to oversell their products. Interestingly, this turns out not to be the case in our model of conflicts for CRAs. The reason is that CRA ratings are as likely to be complements as substitutes and issuers may choose to purchase ratings from both CRAs in equilibrium. Also, the presence of trusting investors distorts CRA incentives to inflate ratings in the same way, whether in a duopoly or a monopoly. In contrast, in Bolton et al. (2007), information revelation comes from the banks' need to differentiate their products.

Several related papers study other implications of shopping for good ratings. Faure-Grimaud, Peyrache, and Quesada (2009) look at corporate governance ratings in a market with truthful CRAs and rational investors. They show that issuers may prefer to suppress their ratings if they are too noisy. They also find that competition between rating agencies can result in less information disclosure. Skreta and Veldkamp (2009) and Sangiorgi et al. (2009) also assume that CRAs truthfully relay their information and demonstrate how noisier information creates more opportunity for shopping by issuers to take advantage of a naive clientele.

Farhi, Lerner, and Tirole (2010) are interested in how certifiers such as rating agencies or academic journals position themselves with respect to the

transparency and coarseness of their certifications. While they allow for heterogeneity among certifiers, they set aside reputation effects and the incentives to produce generous ratings or certifications. They examine the strategy of sellers (our issuers) when they face certifiers that differ in their standards. When a fail for the high-level certification is not disclosed, sellers may opt for an ambitious certification strategy (approaching certifiers with higher standards first) provided the nondisclosure of the fail is not transparent. This strategy is related to ratings shopping, as the result in both cases is that the market does not observe negative information.

The paper is organized as follows. In Section II, we develop the model and solve the case for a single CRA. In Section III, we analyze the case of competition between two CRAs. Section IV compares efficiency in the two market structures. Section V takes the conclusion from Section IV that competition decreases efficiency and examines its robustness. Section VI investigates different plans to regulate the credit rating industry. Section VII lays out empirical implications of the model and surveys the evidence. Finally, Section VIII concludes.

I. The Model

We consider three types of *risk-neutral* agents: issuers, CRAs, and investors with a measure of one. Funds from investors are sought by issuers for independent investments in multiple periods, although we focus primarily on the analysis of a single issue in the first period.

An investment is characterized by its probability of default: a *bad* investment defaults with probability $p > 0$, and a *good* investment defaults with probability zero. Either type of investment yields the same return R when not in default, and zero in default.[12] The investment has constant returns to scale, so that each unit issued has the same return profile.

All agents believe ex ante that the investment is good with probability $\frac{1}{2}$. This creates a role for the CRA, which can use its technology to find out whether the investment is good or bad. A signal $\theta \in \{g, b\}$, which is the private information of the CRA, has the following informational content about the true type ω of the investment:

$$\Pr(\theta = g \mid \omega = g) = \Pr(\theta = b \mid \omega = b) = e.$$

The variable e measures the quality of the signal received, which we refer to as the *precision* of the signal. At $e = \frac{1}{2}$ the signal has revealed no information and agents retain their ex ante beliefs. For $e > \frac{1}{2}$, the signal is informative. We assume that the level of precision is known and lies in the interval $(\frac{1}{2}, 1)$.

The CRAs post their fee ϕ at which a rating can be purchased before they receive the signal. When they are approached by an issuer, CRAs proceed to

[12] In the working paper version of this model, we allowed for a positive recovery value conditional on default. All of the same results continue to hold, so we have chosen this specification for expositional purposes.

retrieve the signal θ and produce a credit report. After observing the report, the issuer chooses either to pay ϕ to have the CRA's proposed rating distributed or to refuse to purchase it. In other words, we allow the issuer to "*shop*" for ratings. This timing is meant to capture in a simple way the back-and-forth negotiations that often go on when CRAs make their ratings reports.[13] If the issuer shops and refuses to buy the CRA's report, that in itself is a signal, which conveys information to investors.

The published rating is a message or report of $m = G$ ("Good") or $m = B$ ("Bad") that is observable to investors. Once the rating is announced, or if it is not announced due to the issuer's refusal to purchase it, the issuer sets a uniform price T for the investment. Since the cost of the investment's production is normalized to zero, we can interpret the price T as a *spread*. After observing the rating and the price T, investors finally decide how much of the investment to purchase.

There are two types of investors, *sophisticated* and *trusting*. A fraction $1 - \alpha$ of investors is sophisticated. These investors observe the payoffs of the game for both the CRA and the issuer, and therefore understand the CRA's and issuer's potential conflict of interest. However, they do not know whether the investment is good or bad, as they do not observe the signal of the CRA and they only have access to the CRA's report. Trusting investors assume that CRAs always truthfully rate the investment and therefore take CRAs' ratings at face value. Also, when they don't observe a rating, these investors simply retain their ex ante beliefs. Sophisticated investors, in contrast, rationally update their beliefs.

One way to motivate the coexistence of trusting and sophisticated investors is to observe that different types of investors have different incentives to perform due diligence. Trusting investors may be managing third party investments and their pay may only depend marginally upon the realized return of the assets they manage[14] whereas sophisticated investors may be investing their personal funds or their pay may be more closely tied to realized returns.

If investors find out that a CRA inflated its rating, they *punish* the CRA in future periods by ignoring its reports. At the time the rating is issued, however, investors cannot determine whether the rating is truthful; more formally, investors cannot determine whether the rating $m \in \{B, G\}$ is equal to the signal received by the CRA $\theta \in \{b, g\}$. But investors are able to find out ex post whether the CRA lied in the event of a default. In practice, it is difficult to determine whether a CRA misled investors even ex post. Still, it is generally easier to

[13] We do not allow for unsolicited ratings. These ratings are rare in practice (see Sangiorgi et al. (2009)). In the Internet Appendix, we analyze the effects of restructuring an investment (e.g., a structured financial product). We find this is likely to decrease market efficiency, as the sole purpose of the restructuring is to offer a better rating to the trusting investor clientele. (An Internet Appendix for this article is available online in the "Supplements and Datasets" section at http://www.afajof.org/supplements.asp.)

[14] Regulation that forces managers to only purchase investments with good ratings could also provide these incentives. Lower incentives to perform due diligence could also be exacerbated by investments that are more complex and difficult to value.

make that determination ex post rather than ex ante. To simplify the analysis, we make the somewhat extreme assumption that investors can perfectly identify whether the CRA lied in the event of a default.[15] Hence, if the CRA receives a signal $\theta = g$ and reports $m = G$, then if the investment fails the CRA will not be punished, as investors can see that it acted in *good faith*. However, a CRA who receives a signal $\theta = b$ and reports $m = G$ will be punished if the project fails.

Reputation costs create an incentive for CRAs to tell the truth, since investors can eventually learn and punish the CRA. We denote the reputation cost by ρ. This is the discounted sum of future CRA profits, which are available when the CRA is not caught lying.[16] To simplify the analysis, we follow Morgan and Stocken (2003), Ottaviani and Sorensen (2006), and Bolton et al. (2007) by assuming that reputation costs are exogenously given. This allows us to explore policy implications in a tractable manner.

Also for tractability, we also assume that the reputation ρ at stake is slightly noisy:

ASSUMPTION 1: *There is a tiny amount of uncertainty on the part of the CRA about the actual value of* ρ, that is, $\rho \in [\tilde{\rho} - \varepsilon, \tilde{\rho} + \varepsilon]$ such that $\varepsilon \to 0$. *This uncertainty is resolved when the CRA receives its signal.*

This assumption restricts the CRA's strategy space since, for any small amount of uncertainty, however small, it will be unable to set fees exactly at levels to make itself indifferent between reports. Thus, this small uncertainty limits the CRA to pure strategies.

Investors can purchase either one unit or two units of the investment. We assume that they have a reservation utility that is increasing in the size of their investment; specifically, they need a return of u on the first unit of their investment and a return of U on the second unit, where $U > u$.[17] One may think of this in several ways: it could be an investor holding her money in cash and needing a larger return to invest all of it, a need for a higher return in order to commit to only one investment vehicle and not diversify, or a form of risk aversion.

We make the following assumptions on the returns on investment:

ASSUMPTION 2: $(1 - p)R > u$.

[15] Formally, we can motivate this assumption by assuming that the recovery value in default is a random variable and, even though the expected value is normalized to zero, the realizations differ depending on the signal θ observed by the CRA ex ante. The economic idea here is that the issuer also gets a noisy signal θ ex ante and takes greater precautions to salvage some recovery value when $\theta = b$ than when $\theta = g$.

[16] This punishment may be more likely in the case of newer financial instruments like structured finance products where demand for the product may dry up. From a broader perspective, the punishment imposed may be from a change in the regulatory environment due to public outcry, such as enforcing liability claims. Finally, although something similar has not occurred in the recent crisis, the downfall of Arthur Andersen represents a severe punishment to a certification intermediary.

[17] The specific form the reservation utility takes could be modeled in multiple ways and give the same results; this form is chosen for simplicity.

ASSUMPTION 3: $(1-(1-e)p)R > U$.

ASSUMPTION 4: $(1-\frac{p}{2})R < U$.

Assumption 2 says that an investor who knows the investment is bad is willing to purchase one unit. Assumption 3 says that an investor with reliable information that the investment is good is willing to purchase two units. The information problem is explicit in Assumption 4: not knowing whether an investment is good or bad (and evaluating the investment with the ex ante beliefs), an investor is not willing to purchase two units. This implies that, if the CRA did not exist, the issuer would not be able to sell two units to any investor since the probability that the issue is bad is too large. The CRA can therefore potentially improve market efficiency by providing information. These assumptions are standard[18] and are necessary to create a value-enhancing role for CRAs through information provision.

To simplify our expressions for payoffs, we introduce the following notation:

$$V^G = (1-(1-e)p)R - U$$

$$V^B = (1-ep)R - u$$

$$V^0 = \left(1-\frac{p}{2}\right)R - u.$$

The terms V^G and V^B represent the marginal value[19] to sophisticated investors when the CRA truthfully reports $m = G$ and $m = B$, respectively. They also represent the marginal value to trusting investors when the CRA reports $m = G$ and $m = B$, whether truthfully or not. The term V^0 is the marginal value to investors who maintain their ex ante beliefs about the value of the investment.

A. *The Ratings Game with a Single CRA*

We begin by examining the game with a monopoly CRA. The timing of moves in this game is as follows:

1. The CRA posts its fee ϕ.
2. The CRA receives the signal and then makes a report of $m = G$ or $m = B$.
3. The issuer observes the report and decides whether to buy and distribute it. The issuer then sets a price T for a unit of the investment.

[18] For example, Mathis et al. (2009) assume that, in the absence of a CRA, the investment will not be purchased, and the CRA can improve market efficiency by providing information about which investments are good.

[19] We define the marginal value V^B with respect to the first unit of investment and its reservation value u because investors will only purchase one unit of a bad investment. We define the marginal value V^G with respect to the second unit of investment and its reservation value U because investors will purchase two units of an investment they believe to be good, and since issuers are assumed to use uniform pricing, the price must be based on the marginal unit.

4. Investors observe the price T and the CRA rating, if there is any, and decide how much of the investment to purchase.
5. The investment return is realized.

When the monopoly CRA receives a signal it must decide what to report. The issuer must decide whether to purchase the report, and subsequently how much to charge investors. Sophisticated investors must infer how good the investment is and formulate their willingness to pay.[20] We solve the game backwards, beginning with the CRA decision of what report to issue after observing the signal.

LEMMA 1: *Given the fee ϕ, the CRA's reporting strategy is as follows:*
For $\phi > ep\rho$, the CRA inflates ratings (always reports G).
For $0 < \phi < ep\rho$, the CRA reports the truth, relaying its signal perfectly.

The proof is in the Internet Appendix.

When the CRA offers a B rating, the issuer responds by not purchasing this rating, as it only decreases investor valuations. The CRA therefore only obtains the fee ϕ when it offers the G rating. There are thus two possible reporting regimes, one in which the CRA inflates the investment quality (when the fee is larger than the expected reputation cost) and one in which the CRA truthfully reveals the investment quality (when the fee is smaller than the expected reputation cost).

We proceed next to derive the equilibrium fees the CRA sets under each informational regime.

PROPOSITION 1: *The equilibrium of the fee setting game is*

1. If $\alpha 2V^G - V^0 > ep\rho$, the CRA inflates ratings, sets $\phi = \alpha 2V^G - V^0$, and has profits

$$\alpha 2V^G - V^0 + \left(1 - \frac{ep}{2}\right)\rho.$$

2. If $\alpha 2V^G - V^0 < ep\rho$, the CRA reports truthfully, sets $\phi = \min[2V^G - \max[\alpha V^0, V^B], ep\rho]$, and has profits

$$\frac{1}{2}\min[2V^G - \max[\alpha V^0, V^B], ep\rho] + \rho.$$

The proof is in the Internet Appendix.

Proposition 1 establishes that the CRA maximizes its profits by choosing ratings inflation over truth telling whenever the profits from ratings inflation ($\alpha 2V^G - V^0$) are larger than the expected reputation cost $ep\rho$.[21] Overstating

[20] There are situations in which a report of "Bad" ($m = B$) is off the equilibrium path. Because we employ the concept of Perfect Bayesian Equilibrium, there is no restriction on off-the-equilibrium-path beliefs. However, we restrict attention to equilibria in which off-equilibrium-path beliefs are equal to ex ante beliefs (i.e., the investment is expected to be *good* with probability $\frac{1}{2}$).

[21] The fee $\alpha 2V^G - V^0$ represents selling two units of the investment to each trusting investor, who believes the G rating. This fee must subtract off V^0, because the issuer must be

the quality of the investment is an equilibrium outcome, despite the presence of reputation costs. This is also a point that Mathis et al. (2009) make.

The cutoff $\alpha 2V^G - V^0 - ep\rho$ determines whether the CRA inflates the quality of the investment. Thus, when reputation costs are smaller and the size of the trusting audience larger, the CRA is more likely to take advantage of trusting investors by inflating ratings. Conversely, when reputation costs are larger and the size of the sophisticated audience larger, the CRA is more likely to tell the truth and create information for all investors. This suggests that ratings inflation is more likely in boom times when investors have lower incentives to perform due diligence, as the ex ante quality of investments is then higher. Note also that an increase in the precision of the signal e has competing effects. It raises the expected valuation of trusting investors, giving higher short-term returns to the CRA. On the other hand, it also increases the likelihood that the CRA gets caught if it misled investors, decreasing future returns.

II. Competition among Ratings Agencies

We next examine the game where two ratings agencies compete in selling ratings to issuers. The CRAs can be thought of as having differentiated products since they are receiving imperfect ($e < 1$) signals about the quality of the investment. In addition, two ratings provide *more* information than just one rating, so the issuer may want to purchase both. The timing of the game with competition is similar to the game with one CRA:

1. Each CRA posts a fee ϕ_k, where $k = 1, 2$ represents the firm.
2. The CRAs receive their signals and produce reports of $m = G$ or $m = B$.
3. The issuer observes the reports and decides whether to purchase and distribute one, both, or neither report. It then sets a price T per unit of the investment.
4. Investors observe the report(s) purchased by the issuer and decide how much of the investment to purchase.
5. The return is realized.

Again to simplify our expressions for payoffs, we adopt the following notation:

$$V^{GG} = \left(1 - \frac{(1-e)^2}{(1-e)^2 + e^2} p\right) R - U,$$

$$V^{BB} = \left(1 - \frac{e^2}{(1-e)^2 + e^2} p\right) R - u.$$

The terms V^{GG} and V^{BB} represent the marginal value to sophisticated investors when both CRAs truthfully report $m = G$ and $m = B$, respectively. They also represent the marginal value to trusting investors when both CRAs report

compensated for deciding to do business with the CRA, rather than sell to investors with ex ante valuations.

$m = G$ and $m = B$, respectively, whether truthfully or not. The marginal value to trusting investors when one CRA reports $m = G$ and the other reports $m = B$ is V^0, the ex ante marginal value before any information about the investment is obtained.

To simplify the analysis, we make the following assumption about the marginal value of an additional positive report:[22]

ASSUMPTION 5: $\alpha 2V^G - V^0 > 2(V^{GG} - V^G)$.

This means that the value of the first G report for trusting investors is larger than the value of a second G report for all investors. This assumption is a way of expressing decreasing returns to G reports. It is slightly stronger than standard decreasing returns. This assumption is always satisfied if the precision e is sufficiently high (close to one) or sufficiently low (close to one-half), as in those cases $V^{GG} = V^G$.

We also make the following assumption:

ASSUMPTION 6: $\alpha 2V^G - V^0 - \min[2(V^{GG} - V^G), ep\rho^D] < e^2 p\rho^D$.

This condition guarantees existence of the truth telling equilibrium by preventing one CRA from unilaterally deviating to inflating its ratings and catering only to trusting customers. If this condition did not hold, there would be less truth telling in duopoly, which would only strengthen our results on the efficiency of monopoly in Sections IV and V.

Finally, we make the following assumption:

ASSUMPTION 7: $\frac{V^0}{2V^G} < \frac{V^{BB}}{V^0}$.

This is not critical to the results at all, but simplifies the exposition.[23]

We denote the discounted sum of future profits in a duopoly for each CRA if it is not caught lying by ρ^D. Again this is an exogenous variable as in the case of monopoly.[24] As before, we solve the game backwards, beginning with the decision of what report to issue after observing the signal.

[22] Without Assumption 5, there can still be equilibria in which both CRAs tell the truth and equilibria in which both CRAs always report G (and there would be no equilibria in which one CRA tells the truth and one always reports G). However, there would be multiple equilibria for each informational regime, there would need to be another restriction on parameters to guarantee existence, and both types of equilibria could coexist. Assumption 5 also places a lower bound on α, which means some shopping will always occur. This plays a role in our analysis of market efficiency.

[23] Assumption 7 fixes the cutoffs for which shopping will occur. When there are two B reports the issuer must decide between charging V^0 to trusting investors or V^{BB} to everyone. There is then a cutoff $\frac{V^{BB}}{V^0}$ for α such that it is best to target trusting investors for α higher than this cutoff (when there are two B reports), that is, $\max[\alpha V^0, V^{BB}] = \alpha V^0$. This will be relevant in the proposition below and in the section on market efficiency.

[24] Note that this is a stronger assumption, since with two CRAs it might be the case that should one CRA be caught lying, the other CRA gets larger continuation profits. It might also be the case that a CRA only gets caught if it is lying and the other CRA is telling the truth (Stolper (2009) examines this type of reputation in a game where a regulator is actively monitoring and punishing CRAs). Alternatively, it might be the case that both CRAs (the whole industry) get punished when any CRA is caught. Furthermore, there may be an added difference between monopoly and duopoly in the sense that when a monopoly CRA gets caught there is nowhere to turn, while when a duopoly

LEMMA 2: *For a given set of fees for both CRAs, CRA k's reporting strategy is as follows:*

1. If $\phi_k > ep\rho^D$*, the CRA inflates ratings (always reports G).*
2. If $\phi_k < ep\rho^D$*, the CRA reports the truth, relaying its signal perfectly.*

The proof is the same as that of Lemma 1. We next solve for the equilibrium of the fee-setting game.

PROPOSITION 2: *The equilibrium of the fee-setting subgame (assuming Assumptions 5 to 7 hold) is as follows:*

1. If $\alpha 2(V^{GG} - V^G) > ep\rho^D$*, both CRAs always report G,* $\phi_k = \alpha 2(V^{GG} - V^G)$ *for* $k = 1, 2$*, and CRA profits are given by*

$$\alpha 2(V^{GG} - V^G) + \left(1 - \frac{ep}{2}\right)\rho^D.$$

2. If $\alpha 2(V^{GG} - V^G) < ep\rho^D$*, both CRAs report truthfully,* $\phi_k = \min[2(V^{GG} - V^G), ep\rho^D]$ *for* $k = 1, 2$*, and CRA profits are given by*

$$\frac{1}{2}\min[2(V^{GG} - V^G), ep\rho^D] + \rho^D.$$

The proof is in the Internet Appendix.

There are thus two possible equilibria: one in which both CRAs always inflate the quality of the investment, and one in which both CRAs truthfully reveal their information about the investment. The cutoff determining which equilibrium prevails is whether the current payoff from inflating ratings $\alpha 2(V^{GG} - V^G)$ is larger or smaller than the expected cost of getting caught $ep\rho^D$.

In general, with a larger fraction of sophisticated investors and a larger reputation cost there will be more truth telling. An increase in the precision of the signal, however, creates a trade-off. The probability of getting caught is rising in the precision, making truth telling more likely. But the current payoff from manipulating ($\alpha 2(V^{GG} - V^G)$) is increasing for low precision levels, meaning that truth telling is less likely. However, in contrast to the case of monopoly, for high precision levels the current payoff is decreasing in the precision, meaning that current and future incentives are aligned in making truth telling more likely.

Comparing the outcome under competition to the case of a monopoly CRA—where the cutoff for truth telling is whether $\alpha 2V^G - V^0 - ep\rho$ is larger than zero or not—we find that, as the marginal value of a positive CRA report is decreasing, the payoff to inflating ratings is larger in a monopoly. Still, it is likely that $\rho > \rho^D$, since the expected loss of business should be larger in monopoly, which may mitigate the increase in fees available to the monopolist. Note, however, that if trusting investors were to overestimate the precision of the

CRA gets caught there is a reasonable alternative (the other CRA). This approach is taken in the context of firms selling goods of varying qualities in Hörner (2002).

CRAs' reports, the incentive to inflate would be very strong irrespective of market structure (current payoffs increase, future costs do not change). Ashcraft and Schuermann (2008, p. (ii)) support the idea of overestimation, noting that "Credit ratings were assigned to subprime MBS with significant error. Even though the rating agencies publicly disclosed their rating criteria for subprime, investors lacked the ability to evaluate the efficacy of these models." Finally, if we define shopping as taking place when there are less than two G signals (Pr(Shopping) $= 1 -$ Pr(two G signals)), we find that shopping increases in duopoly when precision decreases.[25] Skreta and Veldkamp (2009) also point out that less precise signals imply more ratings shopping by issuers.

III. Market Efficiency

We now turn to the evaluation of the efficiency of equilibrium outcomes. Note that in our model it is not completely obvious what the relevant efficiency benchmark is, as a fraction of investors are trusting. We consider total ex ante surplus,[26] evaluating expected surplus for all agents from the point of view of a sophisticated agent, thus adopting a paternalistic point of view. In other words, we take the view that one role of financial regulation is to protect trusting investors from mistakes they may make based on faulty information. The main motivation for this view is that trusting investors would support such regulations with the benefit of hindsight once their naivete is exposed.

We begin by establishing two benchmarks for total surplus, the first-best and the market solution when there are no CRAs. The first-best (subscript FB) is given by

$$\begin{aligned} W_{\text{FB}} &= \frac{1}{2}(2R - u - U) + \frac{1}{2}((1-p)R - u) \\ &= V^0 + \frac{1}{2}(R - U). \end{aligned}$$

This expression is given by the probability that the investment is good multiplied by the surplus created when investors purchase two units plus the probability that the investment is bad multiplied by the surplus when only one unit is purchased.

The market solution when there are no CRAs (subscript 0) is given simply by

$$W_0 = V^0,$$

[25] As Charles Calomiris has argued: "Subprime was a relatively new product, [. . .] Given the recent origins of the subprime market which postdates the last housing cycle downturn in the U.S. (1989-1991), how were the rating agencies able to ascertain what the LGD would be on a subprime mortgage pool?" Thus, the lower precision of CRAs' information about subprime credit risk may have been a source of ratings inflation through greater shopping pressure by issuers. "The Subprime Turmoil: What's Old, What's New, and What's Next," by Charles Calomiris, Vox: //www.voxeu.org/index.php?q=node/1561, 2008.

[26] In a previous version of the paper, we also used investor surplus to evaluate market efficiency. The results were the same when comparing truthtelling regimes, but stronger when comparing ratings inflation regimes (duopoly had strictly lower investor surplus than monopoly).

since both trusting and sophisticated investors would then only purchase one unit. The maximum surplus that can be gained through the provision of credit ratings is therefore given by $\frac{1}{2}(R-U)$, the extra unit purchased when the investment is good.

We now analyze the total surplus in each regime for both monopoly and duopoly. In the total surplus calculations we add the surplus of investors, CRAs, and issuers. The fees of CRAs and the prices charged by issuers net out. Note also that we exclude future surplus from our welfare calculations and look only at efficiency in the short run, as our reputation parameters ρ and ρ^D are exogenous. Finally, note that Assumption 5 implies that $\alpha 2V^G - V^0 > 0$, or $\alpha > \frac{V^0}{2V^G}$. We therefore examine total surplus (and investor surplus) only for the interval $\alpha \in [\frac{V^0}{2V^G}, 1]$.

1. Monopoly CRA, ratings inflation regime ($\alpha 2V^G - V^0 > ep\rho$):
   Only trusting investors purchase at the high prices, as the rating reveals no positive information to sophisticated investors. Since trusting investors believe the investment is good, they invest two units. Total surplus is then

$$W_M^G = \alpha[V^0 + (V^0 + u - U)] \tag{1}$$

   (where the subscript M refers to the monopoly and the superscript G refers to the fact that the CRA always reports G).
   This expression is positive, although it may be quite small. The first term in the expression in square brackets is our market solution when there are no CRAs and is positive, while the second term is negative by Assumption 4. Hence, as intuition suggests, the presence of a CRA reduces surplus in this scenario.
2. Monopoly CRA, truth telling regime ($\alpha 2V^G - V^0 < ep\rho$):
   There are two subcases here, depending on how the issuer prices the investment when there is no report (interpreted correctly by sophisticated investors as a B report that was not purchased). First, when $\alpha V^0 < V^B$, the issuer optimally sets its price low enough to sell the issue to both types of investors. Total surplus then equals:

$$W_{MT1} = V^0 + \frac{1}{2}V^G. \tag{2}$$

   As expected, the surplus is higher than when there is no CRA as $V^G > 0$ by Assumption 3. As the precision approaches $e = 1$, the surplus approaches the first-best.
   When $\alpha V^0 > V^B$, there is an additional distortion because the issuer then sets its price high (when there is no report) to cater only to trusting investors. In this subcase, total surplus is obviously smaller:

$$W_{MT2} = V^0 + \frac{1}{2}V^G - \frac{1-\alpha}{2}V^B. \tag{3}$$

3. Duopoly, ratings inflation regime ($\alpha 2(V^{GG} - V^{G}) > ep\rho^{D}$):
   Total surplus here is exactly the same as when there is a monopoly CRA who always reports G. Trusting investors purchase two units and sophisticated investors purchase nothing. The split of rents between CRAs and the issuer, however, is different here, as the issuer can earn more than V^0 per investor due to competition. If there is a fixed operating cost for CRAs, this would be less efficient than the case of a monopoly CRA. Both an inflating monopoly and an inflating duopoly are less efficient than a market without CRAs.
4. Duopoly, truth telling regime ($\alpha 2(V^{GG} - V^{G}) < ep\rho^{D}$):
   When $\alpha \in [\frac{V^0}{2V^G}, \frac{V^{BB}}{V^0}]$, the issuer sets the price so as to cater to both types of investors when there is no report. Total surplus then equals

$$\begin{aligned} W_{DT1} &= (e^2 + 2e(1-e)\alpha + (1-e)^2)V^0 + \frac{1}{2}(e^2 - (1-e)^2)(R-U) \\ &\quad + (2e(1-e)\alpha + (1-e)^2)(V^0 + u - U). \end{aligned} \tag{4}$$

   In contrast, when the fraction of trusting investors is large ($\alpha \in [\frac{V^{BB}}{V^0}, 1]$) and when there are no G reports the issuer sets a high price at which only trusting investors purchase. Trusting investors are also the only ones to purchase when there is only one G report. Thus, the total surplus is the same as in equation (4), minus the surplus lost from the fact that the issuer targets only trusting investors:

$$\begin{aligned} W_{DT2} &= (e^2 + 2e(1-e)\alpha + (1-e)^2)V^0 + \frac{1}{2}(e^2 - (1-e)^2)(R-U) \\ &\quad + (2e(1-e)\alpha + (1-e)^2)(V^0 + u - U) \\ &\quad - \frac{1-\alpha}{2}\left[(1-e)^2(R-u) + e^2((1-p)R - u)\right]. \end{aligned} \tag{5}$$

Comparing these expressions for total surplus in equations (1) to (5), we find a surprising result: Truth telling in duopoly yields a lower surplus than truth telling in monopoly. We establish this result in the following proposition:

PROPOSITION 3: *Given Assumptions 1 to 7, a truth telling monopoly strictly dominates a truth telling duopoly.*

The proof is in the Internet Appendix.

A duopoly is less efficient because there are more opportunities for the issuer to take advantage of trusting investors. This can occur when one CRA reports G and one reports B, or when both report B. In contrast, under a monopoly CRA there is an opportunity to shop only when the monopoly CRA reports B. As a result, issuers set high prices that exclude sophisticated investors from the market when, from an efficiency perspective, they should be participating. Also, the additional information of the second report is wasted. This is predicated on the fact that Assumption 5 places a lower bound on the number of trusting investors, since, clearly, shopping doesn't occur when all investors are sophisticated.

This result, taken together with the fact that a duopoly is as efficient as a monopoly when both are inflating the quality of the investment (and less efficient if we consider operating costs), suggests that competition among information intermediaries may be detrimental when shopping is allowed. More formally, conditional on being in the same informational regime, monopoly increases total surplus. Therefore, policies encouraging entry may not be the best methods to increase efficiency. This result is in line with the evidence presented in Becker and Milbourn (2011), who document less accurate ratings in the corporate bond market due to more competition from Fitch.

A. *Are There Any Benefits to Competition?*

Our result that competition among CRAs reduces market efficiency is obtained by comparing outcomes under a monopoly and a duopoly for the same informational regime. A natural question that arises is whether there can be any benefits to competition when the informational regime differs across market structures. Consider first the comparison of monopoly and duopoly under different information regimes. It is easy to see that a truth telling monopoly not only dominates a truth telling duopoly but also a duopoly in which CRAs inflate ratings. But does a monopoly CRA that inflates ratings dominate a truth telling duopoly? The next lemma establishes that this is *not* the case.

LEMMA 3: *Total surplus for a truth telling duopoly is larger than that of a monopoly CRA that inflates ratings.*

The proof is in the Internet Appendix.

This lemma underscores the harmful effects of CRA ratings inflation relative to issuer shopping. The parameters for which both scenarios can occur simultaneously depend on the intersection of the following two inequalities:

(1) $\alpha 2(V^{GG} - V^{G}) < ep\rho^{D}$, which guarantees that CRAs in a duopoly prefer to rate truthfully, and
(2) $\alpha 2V^{G} - V^{0} > ep\rho$, which ensures that a monopoly CRA prefers to inflate ratings to attract more issuers.

A few observations are worth pointing out. First, these inequalities can only hold in both market structures if the measure α of trusting investors is small. Otherwise, the financial rewards for CRAs from inflating their ratings and overselling the issue to trusting investors are just too high. Second, truth telling in the duopoly is more likely when the informational value of a second rating is low (V^{GG} close to V^{G}). Thus, somewhat paradoxically, a CRA duopoly dominates a monopoly only in situations in which the marginal value of a second CRA is small. Moreover, in the following section we demonstrate that when reputation is endogenized, the CRAs will be in the same information regime for all parameters, which implies that monopoly is always more efficient.

Finally, note that, even when a duopoly dominates a monopoly, this does not imply that competition is efficient, as the negative effects from issuer shopping remain. It is straightforward to show that reducing competition to create a

regulated duopoly, in which issuers are required to purchase a rating from both CRAs, would be welfare superior to an unregulated duopoly.[27] Indeed, under such a regulation (1) CRAs would always strictly prefer to rate truthfully, as the purchase of their rating is then no longer contingent on its content; (2) issuer shopping would be eliminated; and (3) issues would be rated based on the maximum available information. In fact, without the CRA conflict of interest and issuer shopping, total surplus would be equal to the first-best (constrained, of course, by the precision of the CRAs' information).[28]

IV. Endogenous Reputation

We now explore the effect of endogenizing reputation costs in a fully specified dynamic model. The endogenous reputation cost from being caught inflating ratings is the cost in forgone future ratings sales. The simplest way of extending our model to allow for such an endogenous reputation cost is to consider a two-period version in which the payoff weight attached to the second period is given by a parameter β (as, e.g., in Laffont and Tirole (1993)), where β may be larger than one. The size of the parameter β then represents the importance of future relative to current profits for the CRAs. Thus, for example, at the onset of an issuance boom future capitalized CRA profits are likely to be large, so that β is large. In contrast, at the end of an issuance boom and at the onset of a recession β is small.

Consider first the situation of a monopoly CRA. The simplification obtained from the two-period formulation is that we can solve the game backwards starting from the second period (taking as given that the CRA has not been caught inflating ratings in the first period). As the second period is the last period, there are no more reputation concerns that discipline the CRA, so that the CRA always inflates its ratings. From Assumption 5 we know $\alpha 2V^G - V^0 > 0$, so that the CRA's optimal policy in the second period is to sell the overrated issue only to trusting investors and thus realize a positive profit of $\alpha 2V^G - V^0$. In period 1, *endogenous reputation costs* from forgone future profits are then given by $\rho = \beta(\alpha 2V^G - V^0)$. With such an endogenous reputation cost, the CRA then inflates ratings in period 1 if and only if

$$(\alpha 2V^G - V^0) > ep\beta(\alpha 2V^G - V^0),$$

or

$$\beta < \frac{1}{ep}.$$

This simple analysis of the dynamic CRA monopoly thus reveals that with endogenous reputation costs a CRA is more likely to engage in ratings inflation

[27] In the model, this means purchasing from two CRAs. In practice, realistically this would imply purchasing from the big three CRAs (Moody's, Standard & Poor's, and Fitch).

[28] Issuers may lose out under this regulation if CRAs remain free to set prices since, as under a monopoly, the entire issuer surplus may be appropriated by the CRAs.

when future profits matter less, for example, towards the end of an issuance boom. This result is consistent with both the theoretical results of Mathis et al. (2009) and the empirical findings of Ashcraft, Goldsmith-Pinkham, and Vickery (2010).

Consider next the situation of a duopoly CRA. Once again, the two CRAs inflate ratings in period 2, as there are no costs in being caught inflating ratings. Each CRA's best response in the second period is to sell the overrated issue only to trusting investors and thus realize a positive profit of $\alpha 2(V^{GG} - V^{G})$. In period 1, then, *endogenous reputation costs* from foregone future profits are given by $\rho^{D} = \beta\alpha 2(V^{GG} - V^{G})$. In period 1 a CRA duopoly that inflates ratings, in which each CRA earns $\alpha 2(V^{GG} - V^{G})$, is then an equilibrium if and only if

$$\alpha 2(V^{GG} - V^{G}) > ep\beta\alpha 2(V^{GG} - V^{G}),$$

or again

$$\beta < \frac{1}{ep}.$$

This implies the following result:

PROPOSITION 4: *In the model with endogenous reputation costs, in the first period*

(1) the CRA(s) report truthfully in both monopoly and duopoly iff $\beta \geq \frac{1}{ep}$, *and*
(2) the CRA(s) inflate ratings in both monopoly and duopoly iff $\beta < \frac{1}{ep}$.

Thus, with endogenous reputation costs, it is the same condition that determines whether a monopoly CRA or a duopoly CRA will rate truthfully in period 1 or not. In other words, in our simple dynamic extension with endogenous reputation costs, the equilibrium information regime is the same across market structures, so that a monopoly always dominates a duopoly in this situation. This simple analysis thus suggests that making reputation endogenous may well strengthen our efficiency results rather than weaken them. It would be of interest (but beyond the scope of this paper) to explore these issues more systematically in a fully general dynamic game, possibly with an infinite horizon. There is currently no model of oligopolistic competition over an infinite horizon in the CRA literature; indeed, there are very few such models in the industrial organization literature for obvious reasons of tractability.[29]

V. Regulating the Credit Ratings Industry

The subprime crisis has brought to light the poor performance of CRAs in rating structured financial products and reminded investors of CRAs' poor past performance in predicting the East Asian crisis, the excesses of the dotcom bubble, and the collapse of Enron. Governmental bodies have been debating how to regulate CRAs, and some initial rules have thus been issued.

[29] See Bar-Isaac and Tadelis (2008) for a review of the literature.

In this section we discuss the most prominent proposals in the context of our model. In our view, the key issues that the proposals seek to address are as follows:

(1) eliminate the CRA conflicts of interest by preventing issuers from influencing ratings,
(2) prevent issuers from shopping for ratings and disclosing only those ratings they prefer, and
(3) monitor the quality of the ratings methodology.

New York State Attorney General Andrew Cuomo reached an agreement[30] with credit ratings firms to change some features of the rating process in the summer of 2008. The agreement between Cuomo and Standard & Poor's, Moody's, and Fitch essentially addresses the first point, preventing issuers from paying for specific ratings and forcing issuers to pay the CRA up front before it conducts its initial analysis.[31] This restriction can eliminate ratings inflation by CRAs in our model by eliminating the issuer's ability to provide incentives for good ratings.[32] Importantly, however, it does not eliminate shopping by an issuer; an issuer may still reach an agreement with a CRA to not publish a bad rating. In our model, issuer shopping can create distortions even with unbiased CRA ratings due to the trusting nature of some investors. There have been several moves to decrease CRA conflicts of interest. The SEC recently enacted a rule that prohibits consulting activity related to ratings by CRAs. The Dodd-Frank financial reform bill further states that ratings must be explicitly separated from sales and marketing.[33]

Prohibiting shopping by requiring that CRAs automatically disclose any rating paid for by an issuer would achieve the first-best surplus[34] in our model when combined with the Cuomo plan. The SEC currently has a proposed rule that would formalize this prohibition. Nevertheless, shopping may be difficult to eliminate because of informal discussions between issuers and CRAs that may still take place. This points to a possible need for auditing by a regulator.

[30] The agreement is reportedly for 3 years and on structured finance products only, see "Big Credit-Rating Firms Agree to Reforms," by Aaron Lucchetti, *Wall Street Journal*, June 6, 2008.

[31] There is a fine point here, namely, the deal specifies up-front payments for initial analysis but does not prevent subsequent payments. This is obviously an issue, but outside the scope of our model.

[32] It is possible, of course, to dynamically create these incentives through repeated interactions between an issuer and a CRA. Such analysis is out of the scope of this paper, but the point is certainly a caveat.

[33] An intriguing proposal to diminish conflicts of interest was taken out of the Dodd-Frank bill at the last minute (and was relegated to be the subject of a formal study by the SEC). The "Franken Amendment" proposed to set up a body that would randomly assign issuers of structured finance products to rating agencies. If fees were paid up front to the body, this could eliminate both conflicts of interest and shopping.

[34] While quite intuitive, we prove this formally in the Internet Appendix. An interesting unintended consequence of eliminating shopping is to reduce the number of ratings, which occurs because the rents to an extra rating decrease. This is not a complete surprise, as our model demonstrates that monopoly is more efficient than duopoly.

Even by eliminating shopping from the Cuomo plan, there is a risk of efficiency loss due to moral hazard. Suppose the precision of the signal e is a choice variable of the CRA and larger precision is more costly. If the CRA can choose this precision after being paid up front[35] and it is noncontractible, then the CRAs would choose the minimum precision of $\frac{1}{2}$ and, knowing this, the issuer would not hire the CRA in the first place. There would thus be a breakdown in the market for certification.

Interestingly, our main model with no regulation shows that adding the observable choice of precision in monopoly will lead to positive investment by the CRA since the issuer pays contingent fees. Still, our total surplus calculations show that breakdown of the CRA market could still be a better outcome than one in which a CRA inflates quality, but worse than an outcome in which a CRA tells the truth. Consequently, it is crucial that the new regulatory structure for CRAs be accompanied by oversight of minimum analytical standards for the CRAs (and that these standards be enforceable), so as to regain the beneficial aspects summarized earlier. The Dodd-Frank bill discusses analytical standards and mandates that the SEC issue rules regarding both training, experience, and competence for CRA analysts and procedures and methodologies for the ratings.

One last approach to improving ratings quality lies in enhancing the market's ability to punish CRAs. In our model, this would increase CRA's reputation cost, making truth telling more likely. The SEC has issued rules forcing CRAs to disclose their track record, making their performance more transparent. More importantly, the Dodd-Frank bill lowers the bar for liability claims against CRAs. CRAs had been "immune from liability for misstatements ... under Section 11 of the Securities Act of 1933" and have won most cases against them based on arguments that credit ratings are free speech and "extensively disclaimed" (Partnoy (2002, pp. 78, 79)). Eliminating this immunity could therefore impose serious costs on CRAs for ratings inflation.

VI. Empirical Implications

This paper demonstrates that competition among CRAs can reduce market efficiency due to the shopping effect and provides a framework for understanding the trade-offs in recent policy proposals regarding the credit rating industry. In this section we examine evidence surrounding testable implications of the model. We conclude by discussing systematic evidence related to our assumption that investors are trusting.

The model offers several testable hypotheses:

1. The model shows that poor quality ratings are increasing in the fraction of trusting investors and current payoffs, and decreasing in the expected probability of getting caught (the reputation cost). While it is difficult to

[35] If precision were not effort in performing the analysis on the investment, but rather quality of the analytic models used, then the CRAs would choose greater precision, since the level of precision would entirely be chosen *before* the fees were paid.

measure these variables, all three of these factors are more likely to occur during boom times and less likely to occur during recessions: when times are good, the probability of defaults is lower, which may decrease due diligence on the part of investors as well as evidence of ratings bias. As a follow-on effect, this can increase demand and issuance, generating larger fees for CRAs. To test this implication of the model, we can examine whether poor ratings quality is more likely during boom times.

2. As opposed to other theoretical papers, ratings inflation can arise in our model purely due to conflicts of interest and not shopping. Empirically, we can attempt to exploit the lack of shopping possibilities in the corporate bond markets to examine whether conflicts of interest (the trade-off between higher current profits and expected future profits) have some explanatory power. Variables that affect current and future profits, such as the degree of competition, can be related to ratings quality.
3. We show in the section on competition that shopping is more likely when the CRAs' models are less precise.
4. Shopping is used by issuers to exploit trusting investors. Hence, if shopping occurs, investors are not taking into account the selection effect. These observations can be used to test two implications of the theory, namely (i) yields are more dependent on ratings when there is more shopping-type behavior, and (ii) fewer published ratings predict worse ex post performance.

We make use of recent empirical papers on CRAs and ratings quality to examine our hypotheses. To do so, we interpret investment in our model broadly as applying to both the corporate bond markets and structured finance products (indeed, in the Internet Appendix we explicitly model the restructuring process). We attempt to point out where institutional details benefit or detract from the model.

The implication that ratings inflation is more likely to happen during booms has been documented in several recent papers. Ashcraft et al. (2010) find that as MBS issuance volume shot up between 2005 and mid-2007, ratings quality declined. Specifically, subordination levels[36] for subprime and Alt-A MBS deals decreased over this period when conditioning on the overall risk of the deal.[37] Moreover, subsequent ratings downgrades for the 2005 to mid-2007 cohorts were dramatically larger than for previous cohorts. Griffin and Tang (2010) find that CRA adjustments to their models' predictions of credit risk in the CDO market were positively related to future downgrades. These adjustments were overwhelmingly positive and the amount adjusted (the width of the AAA tranche) increased sharply from 2003 to 2007 (from 6% to 18.2%). The adjustments are not well explained by natural covariates (such as past deals

[36] The subordination level they use is the fraction of the deal that is junior to the AAA tranche. A smaller fraction means that the AAA tranche is less "protected" from defaults, and therefore less costly from the issuer's point of view.

[37] Alt-A mortgages, short for Alternative A-paper, are those mortgages that have prime borrowers but a nonstandard characteristic, e.g. regarding documentation. This makes them riskier than prime mortgages but less risky than subprime ones.

by collateral manager, credit enhancements, and other modeling techniques). Furthermore, 98.6% of the AAA tranches in their sample CDOs failed to meet the CRAs' reported AAA standard (for their sample from 1997 to 2007). They also find that adjustments increase CDO value on average by $12.58 million per CDO.

On the relationship between current payoffs and ratings inflation, He et al. (2010) find that MBS tranches sold by larger issuers[38] performed significantly worse (market prices decreased) than those sold by small issuers during the boom period of 2004 to 2006. Faltin-Traeger (2009) shows that when one CRA rates more deals for an issuer in a half-year period than another CRA, the first CRA is less likely to be the first to downgrade that issuer's securities in the next half-year. He also finds that, if a CRA rates a deal higher, that CRA is more likely to be chosen by the issuer on the issuer's next deal. This effect is strongest for Fitch.

Our model isolates two basic causes of poor ratings quality: conflicts of interest (ratings inflation) and shopping. While it is difficult to isolate these in reality, an interesting comparison arises between the corporate bond market and the structured finance market. First, in the corporate bond market, Standard & Poor's and Moody's rate virtually every rated issue. This implies that there is little scope for shopping in that market. Nevertheless, our model suggests that the trade-off between current profits and future payoff may still influence ratings quality. Becker and Milbourn (2011) find supporting evidence: they show that increases in market share by Fitch (a proxy for more competition) lead to higher ratings. Moreover, this evidence suggests that more competition may not be better, even when shopping is not as much of an issue.

Second, the methodology for rating corporate bonds is more standardized and the bonds themselves are much less complicated than structured finance products. Our paper suggests that shopping is more likely when the CRAs' models are less precise, which is certainly the case comparing corporate bonds to structured finance. Within the structured finance arena, Ashcraft et al. (2010) find that the MBS deals that were most likely to underperform were the ones with more interest-only loans (because of limited performance history) and lower documentation, that is, loans that were more opaque or difficult to evaluate.

Our paper posits that shopping is used by issuers to exploit trusting investors. Regarding the dependency of yields on ratings, Adelino (2009) shows that AAA tranche yields of MBS do not have extra predictive power about defaults or subsequent rating downgrades outside of the rating itself. However, it is not obvious from his results that this got worse during the boom (Table 12, Adelino (2009)). With respect to lower ratings leading to poorer performance, there is mixed evidence. Griffin and Tang (2010) find no evidence that CDOs rated by multiple rating agencies experience less default. Both Benmelech and Dlugosz (2009) (for asset-based securities (ABS)) and Ashcraft et al. (2010) (for RMBS), however, find that ex post downgrades of structured

[38] They define larger by market share in terms of deals. As a robustness check, they also look at market share in terms of dollars and find similar results.

finance products are both more likely and larger in deals rated by a single CRA.[39] In preliminary work, Ashcraft et al. (2009) find that the more issuers switch among CRAs, the lower is subordination for Alt-A RMBS, indicating benefits to shopping around.[40]

While not a prediction, a key part of the paper is our assumption that a fraction of investors are trusting. While there is substantial anecdotal evidence to support this assumption, we take this opportunity to describe systematic evidence. There are two views of trusting investors to explain why they do not perform proper due diligence and analysis. The first explains such belief using incentive problems, whereas the second claims that the analysis is too complex for them. While the second is difficult to measure, two papers in the literature on corporate bond ratings demonstrate that ratings are important to some investors solely for regulatory purposes. Kisgen and Strahan (2010) demonstrate that Dominion Bond Rating Service's acquisition of NRSRO status in 2003 changed the impact of its ratings on bond yields only in situations in which this status was important.[41] Bongaerts, Cremers, and Goetzmann (2011) find that Fitch's ratings are often used to break ties between S&P and Moody's.

Focusing on the structured finance market, Adelino (2009) finds intriguing evidence of naivete. He finds that, while initial yields on tranches below AAA predict future credit performance (probability of default and future ratings downgrades), the initial yields on AAA tranches have no predictive power. This is consistent with the hypothesis that investors in AAA tranches have no other information beyond the credit ratings themselves.

VII. Conclusion

Our paper analyzes CRAs and their conflicts of interest. The model includes the critical elements of the industry: Issuer's payments may influence ratings, issuers may shop for ratings, CRA models may vary in precision, barriers to entry create market power for CRAs, reputation considerations affect decision making, and different clienteles for investments exist. These elements allow us to provide a simple general framework for analyzing the rating industry and its efficiency.

Our model yields a surprising result on the adverse effects of competition. In particular, we find that the presence of more trusting investors or lower reputation costs gives CRAs incentives to inflate investment quality, while the precision of the CRAs analysis has dual effects: More precision raises current payoffs but also increases the probability of paying a reputation cost. Our analysis of market efficiency makes it clear that, in general, a monopoly is more efficient than a duopoly. This is because a duopoly provides more opportunities for the issuer to shop and mislead trusting investors. In terms of regulation,

[39] Ashcraft et al. (2009) note, however, that only 1% of their deal sample has just one rating.

[40] The sign is the same for subprime deals but the coefficient is not significant.

[41] That is, regulations that required investments to use the best or second-best NRSRO rating and specifically around the investment grade threshold.

we suggest that up-front fees (as in the Cuomo plan) accompanied by automatic disclosure of ratings and oversight of analytical standards will minimize distortions from conflicts of interest and shopping.

To present a closed-form model of CRA competition, we abstract away from several aspects of the industry that would be worth analyzing further. We simplify the ratings process to allow for only two levels of ratings rather than a finer partition and do not allow for subsequent upgrades and downgrades that CRAs make while monitoring an investment. In terms of the investment being issued, we do not model conduits with multiple assets or make a clear distinction between idiosyncratic and systematic risk. Finally, while we did extend the model in Section V to two periods to endogenize reputation, a model of CRA competition over a longer time horizon could yield interesting results.

REFERENCES

Adelino, Manuel, 2009, Do investors rely only on ratings? The case of mortgage-backed securities, Working paper, Dartmouth.

Ashcraft, Adam, Paul Goldsmith-Pinkham, and James Vickery, 2009, The role of incentives and reputation in the rating of mortgage-backed securities, Working paper, Federal Reserve Bank of New York.

Ashcraft, Adam, Paul Goldsmith-Pinkham, and James Vickery, 2010, MBS ratings and the mortgage credit boom, Working paper, Federal Reserve Bank of New York.

Ashcraft, Adam, and Til Schuermann, 2008, Understanding the securitization of subprime mortgage credit, Working paper, Federal Reserve Bank of New York.

Bar Isaac, Heski, and Joel Shapiro, 2010, Ratings quality over the business cycle, Working paper, NYU.

Bar Isaac, Heski, and Steve Tadelis, 2008, Seller reputation, *Foundations and Trends in Microeconomics* 4, 273–351.

Becker, Bo, and Todd Milbourn, 2011, How did increased competition affect credit ratings? *Journal of Financial Economics* 101, 493–514.

Benabou, Roland, and Guy Laroque, 1992, Using privileged information to manipulate markets: Insiders, gurus, and credibility, *Quarterly Journal of Economics* 107, 921–958.

Benmelech, Efraim, and Jennifer Dlugosz, 2009, The credit rating crisis, *NBER Macroeconomics Annual*, 161–207.

Biglaiser, Gary, 1993, Middlemen as experts, *RAND Journal of Economics* 24, 212–223.

Bolton, Patrick, Xavier Freixas, and Joel Shapiro, 2007, Conflicts of interest, information provision, and competition in the financial services industry, *Journal of Financial Economics* 85, 297–330.

Bongaerts, Dion, K.J. Martijn Cremers, and William N. Goetzmann, 2011, Tiebreaker: Certification and Multiple Credit Ratings, *Journal of Finance*, forthcoming.

Boot, Arnoud W.A., Todd T. Milbourn, and Anjolein Schmeits, 2006, Credit ratings as coordination mechanisms, *Review of Financial Studies* 19, 81–118.

Durbin, Erik, and Ganesh Iyer, 2009, Corruptible advice, *American Economic Journal: Microeconomics* 1, 220–242.

Faltin-Traeger, Oliver, 2009, Picking the right rating agency: Issuer choice in the ABS Market, Working paper, Columbia Business School.

Farhi, Emmanuel, Josh Lerner, and Jean Tirole, 2010, Fear of rejection? Tiered certification and transparency, Working paper, Harvard University.

Faure-Grimaud, Antoine, Eloic Peyrache, and Lucia Quesada, 2009, The ownership of ratings, *Rand Journal of Economics* 40, 234–257.

Griffin, John M., and Dragon Y. Tang, 2010, Did subjectivity play a role in CDO credit ratings? Working paper, UT-Austin.

Harrison, J. Michael, and David M. Kreps, 1978, Speculative investor behavior in a stock market with heterogeneous expectations, *Quarterly Journal of Economics* 92, 323–336.

He, Jie, Jun Qian, and Philip E. Strahan, 2010, Credit ratings and the evolution of the mortgage-backed securities market, Working paper, Boston College.

Hirshleifer, David, and Siew Hong Teoh, 2003, Limited attention, information disclosure and financial reporting, *Journal of Accounting and Economics* 36, 337–386.

Hörner, Johannes, 2002, Reputation and competition, *American Economic Review* 92, 644–663.

Inderst, Roman, and Marco Ottaviani, 2009, Misselling through agents, *American Economic Review* 99, 883–908.

Kartik, Navin, 2009, Strategic communication with lying costs, *Review of Economic Studies* 76, 1359–1395.

Kartik, Navin,, Marco Ottaviani, and Francesco Squintani, 2007, Credulity, lies, and costly talk, *Journal of Economic Theory* 134, 93–116.

Kisgen, Darren J., and Philip E. Strahan, 2010, Do regulations based on credit ratings affect a firm's cost of capital? *Review of Financial Studies* 23, 4324–4347.

Laffont, Jean-Jacques, and Jean Tirole, 1993, *A Theory of Incentives in Procurement and Regulation* (MIT Press, Cambridge).

Lizzeri, Alessandro, 1999, Information revelation and certification intermediaries, *RAND Journal of Economics* 30, 214–231.

Mariano, Beatriz, 2008, Do reputational concerns lead to reliable ratings, Working paper, Universidad Carlos III.

Mathis, Jerome, Jamie McAndrews, and Jean-Charles Rochet, 2009, Rating the raters: Are reputational concerns powerful enough to discipline rating agencies? *Journal of Monetary Economics* 56, 657–674.

Morgan, John, and Philip C. Stocken, 2003, An analysis of stock recommendations, *RAND Journal of Economics* 34, 183–203.

Ottaviani, Marco, and Peter Norman Sorensen, 2006, Reputational cheap talk, *RAND Journal of Economics* 37, 155–175.

Pagano, Marco, and Paolo Volpin, 2010, Securitization, transparency, and liquidity, Working paper, Università di Napoli Federico II and London Business School.

Partnoy, Frank, 2002, The paradox of credit ratings, in Richard M. Levich, Giovanni Majnoni, and Carmen Reinhart, eds.: *Ratings, Rating Agencies and the Global Financial System* (Kluwer Academic Publishers, Boston).

Sangiorgi, Francesco, Jonathan Sokobin, and Chester Spatt, 2009, Credit-rating shopping, selection and equilibrium structure of ratings, Working paper, Carnegie Mellon University.

Scheinkman, Jose, and Wei Xiong, 2003, Overconfidence and speculative bubbles, *Journal of Political Economy* 111, 1183–1219.

Securities and Exchange Commission (SEC), 2008, Summary Report of Issues Identified in the Commission Staff's Examinations of Select Credit Rating Agencies. Available at http://sec.gov/news/studies/2008craexamination070808.pdf.

Skreta, Vasiliki, and Laura Veldkamp, 2009, Ratings shopping and asset complexity: A theory of ratings inflation, *Journal of Monetary Economics* 56, 678–695.

Stolper, Anno, 2009, Regulation of credit rating agencies, *Journal of Banking & Finance* 33, 1266–1273.

Strausz, Roland, 2005, Honest certification and the threat of capture, *International Journal of Industrial Organization* 23, 45–62.

White, Lawrence J., 2002, The credit rating industry: An industrial organization analysis, in Richard M. Levich, Giovanni Majnoni, and Carmen Reinhart, eds.: *Ratings, Rating Agencies and the Global Financial System* (Kluwer Academic Publishers, Boston).

White, Lawrence J., 2010, Markets: The credit rating agencies, *Journal of Economic Perspectives* 24, 211–226.