

**AN EXPERIMENTAL INVESTIGATION
OF RATING-MARKET REGULATION**

Claudia Keser, Asri Özgümüs, Emmanuel Peterlé,
Martin Schmidt

GEORG-AUGUST-UNIVERSITÄT GÖTTINGEN

An experimental investigation of rating-market regulation

Claudia Keser^{a,*}, Asri Özgümüs^a, Emmanuel Peterlé^b, Martin Schmidt^c

^aUniversity of Göttingen, e-mail: claudia.keser@uni-goettingen.de;
asri.oezguemues@wiwi.uni-goettingen.de

^bUniversité de Franche-Comté, e-mail: emmanuel.peterle@univ-fcomte.fr

^cKarlsruhe Institute of Technology (KIT), e-mail: martin.schmidt@kit.edu

December 08, 2016

Abstract: We introduce a simple game-theoretical model that captures the main aspects of the repeated interaction between an issuer and a credit-rating agency. It involves up-front payments of issuer-fees and direct publication of requested ratings. Due to pecuniary injuries for untruthful ratings, the credit-rating agency should always report truthfully in the subgame perfect equilibrium. Knowing this, the issuer should never request a rating. Conducting laboratory experiments, we find that behavior significantly deviates from the equilibrium prediction in favor of a cooperative solution: issuers frequently do request ratings, which is often reciprocated with untruthful good ratings.

Keywords: Game theory, laboratory experiments, rating agencies, regulation.

JEL Codes: C70, C9, G0

* *Corresponding author:* Chair of Microeconomics and Göttingen Laboratory of Behavioral Economics, Georg-August-Universität Göttingen, Platz der Göttinger Sieben 3, D - 37073 Göttingen / Germany

1. Introduction

Credit-rating agencies (CRAs) play an important role in generating and disseminating information about the creditworthiness of a firm, government or bank, and in estimating a financial instruments' likelihood of default. The role of CRAs, however, has been questioned since highly rated structured finance products defaulted en masse during the U.S. subprime mortgage crisis in 2007/8 to 2009. According to the Financial Crisis Inquiry Commission (2011), the financial crisis would not have occurred if today's most famous CRAs Moody's, Standard&Poor's (S&P), and Fitch had performed duly. The poor performance of "The Big Three" in providing timely and accurate ratings can be seen as the result of conflicts of interest in the credit rating market, which have their roots in the history of CRAs, particularly in the early 1970s.

In October 1970, Moody's modified its compensation scheme from "investor-pays" to "issuer-pays", followed by S&P in July 1974.¹ While this change was intended to prevent investors from free riding, it initiated other conflicts of interest in the rating market. The issuer-pays model is characterized by two main difficulties. First, a rating agency is paid only if an issuer asks for the publication of the rating; that is, issuers are able to observe the type of rating before making the buying decision.² Second, CRAs receive over 90% of their revenues from issuer fees (e.g., Darcy 2009), creating a high dependency of CRAs on those fees and therefore, a strong bargaining power for (big) issuers.

Recent studies show that the issuer-pays model makes rating agencies offer (overly) optimistic ratings in order to attract business (e.g., White, 2010; Strobl and Xia, 2011; Jiang et al., 2012; Camanho et al., 2012; Bongaerts, 2014), a phenomenon which is also known as "ratings inflation". Furthermore, the issuer-pays model leads to "ratings shopping" in that it allows issuers to

¹ E.g., Jiang et al. (2012).

² E.g., Securities and Exchange Commission (2008).

solicit multiple ratings and choose the most favorable one (e.g., Skreta and Veldkamp, 2009; Deb and Murphy, 2009; White, 2010; Bolton et al., 2012).

A first attempt to regulate the issuer-pays compensation model was made in 2008 by Andrew Cuomo, then Attorney General and today's governor of New York. Cuomo reached a three-year agreement with Moody's, S&P, and Fitch, which required issuers of structured finance products to pay the issuer-fee up front, before a CRA conducts its initial analysis.³ The reform aimed for an increase in the overall accuracy of ratings by making CRAs less dependent on the question of receiving the rating-fee or not.

The reform, however, is not free from critique. Bolton et al. (2012) present a theoretical rating-market model demonstrating the conflicts of interest between CRAs and issuers.⁴ In the context of their model they argue that Cuomo's agreement does not prevent issuers from ratings shopping. An issuer might consider up-front payments as sunk costs and thus be willing to pay, ex post, for an agreement with the CRA not to publish a bad rating. They claim that the first-best solution in their model would be a reform that prohibits rating shopping by enforcing CRAs to automatically disclose any rating paid for by an issuer.

While this solution is the result of logical, theoretical reasoning, we might expect the real behavior of agents to be different from the theoretical prediction, in particular when interaction is repeated. In highly concentrated markets, issuers and CRAs are likely to have an interest in a long-term

³ New York State Office of the Attorney General (2008).

⁴ The model involves an incentive for CRAs to provide overly optimistic ratings, which negatively correlates with expected reputation costs for untruthful ratings. Comparing a competitive to a monopolistic market the authors find, in keeping with the empirical observation by Becker and Milbourn (2011), that competition among CRAs potentially lowers the overall quality of ratings. According to Skreta and Veldkamp (2009), this is due to the opportunity for issuers to shop for ratings.

cooperation: issuers are interested in good ratings, which CRAs might be willing to supply since they are interested in future rating requests.⁵ If this is the case, Cuomo's regulatory reform is unlikely to prevent CRAs from inflating ratings, even when rating shopping is impossible and rating disclosure is enforced.

The goal of our study is to throw light on this controversy by conducting a laboratory experiment. We design a simple game-theoretical model, inspired by the monopolistic rating-market model in Bolton et al. (2012). Our issuer-CRA rating game includes up-front payments of issuer-fees (Cuomo's agreement) and, as suggested by Bolton et al. (2012), direct publication of any rating that was requested by an issuer. In addition, we include (low and high) pecuniary injuries for inaccurate ratings, where any inaccuracy in ratings is perfectly assignable to the conscious decision of the CRA. We are not explicit about the nature of these costs; they could, for example, reflect reputational losses among investors or financial penalties.⁶ The basic game is parameterized such that the subgame perfect equilibrium strategy for a CRA is to provide truthful ratings. Knowing this, the issuer should not request a rating. Conducting laboratory experiments on a repeated version of this game, we find that, irrespective of the level of the CRA's pecuniary injuries for false ratings, ratings are requested. Furthermore, we frequently observe CRAs to be willing to falsely provide good ratings, except for the final period of the game.

⁵ Frenkel (2015) points at recent empirical studies by Ashcraft et al. (2010) and Griffin and Tan (2012) that provide evidence of rating inflation in markets for structured assets, such as mortgage-backed securities (MBSs) and collateralized debt obligations (CDOs). These markets are highly concentrated. Frenkel (2015) argues that CRAs have an incentive to build not only a public reputation among investors but also a second, private reputation for "leniency" or, in our words, a willingness to cooperate among issuers in order to attract future business. He demonstrates in a theoretical model that in markets with a small number of issuers this may lead to inflated ratings.

⁶ Our model relates to the highly concentrated market model investigated by Frenkel (2015). The focus of our model is on the CRA's reputation building process among issuers, while the concern for public reputation for credibility among investors is controlled for in the parameterization of the CRA's low or high pecuniary injuries for inaccurate ratings.

This leads us to the interpretation that both issuers and CRAs tend to show an interest in a long-term cooperation, which is of pure strategic nature (even though it is different from the subgame perfect equilibrium solution).

In the following Section 2 we present the experimental design. Section 3 provides the experimental results. Section 4 concludes the paper.

2. Experimental Design

2.1 The Game

We consider a repeated two-player game with asymmetric information, in which an issuer- and a CRA-player interact over 20 periods. In each period, the issuer holds a (financial) product, whose quality is contingent upon the state of nature. The quality is either good or bad, each with probability one half.⁷ The quality becomes publicly observable at the end of the period.

Each period is structured as follows. First, nature draws the quality of the product. Then, without knowledge of the drawn quality, the issuer decides whether or not to request a rating from the CRA. If the issuer decides to request a rating, the CRA privately observes the quality and decides which quality to report (whether or not to report the true quality). After the CRA's rating or the issuer's decision not to request a rating, the period is over. The true quality of the product as well as the rating (if requested) is reported to both players before the new period starts.

The payoffs in each period are determined as follows. If the issuer requests a rating from the CRA, a good rating yields a payoff of 160 (reflecting a high

⁷ For purposes of comparability between treatments, we drew a random sequence of 20 observations for the true quality and fixed it for all treatments (10100100010100101110), where 1 = good and 0 = bad.

product valuation) and a bad rating yields a payoff of 20 (reflecting a low product valuation) for the issuer. The payoff of the CRA depends on whether the rating is truthful or not: a truthful rating yields a payoff of 100, while the payoff of an untruthful rating is given by a lower amount, which can be presented as 100 minus a pecuniary injury P . The pecuniary injury, which does not explicitly appear as a cost in the experiment instructions, reflects, for example, a reputational loss after an untruthful rating has become public or regulatory sanctioning costs. If the issuer does not request a rating from the

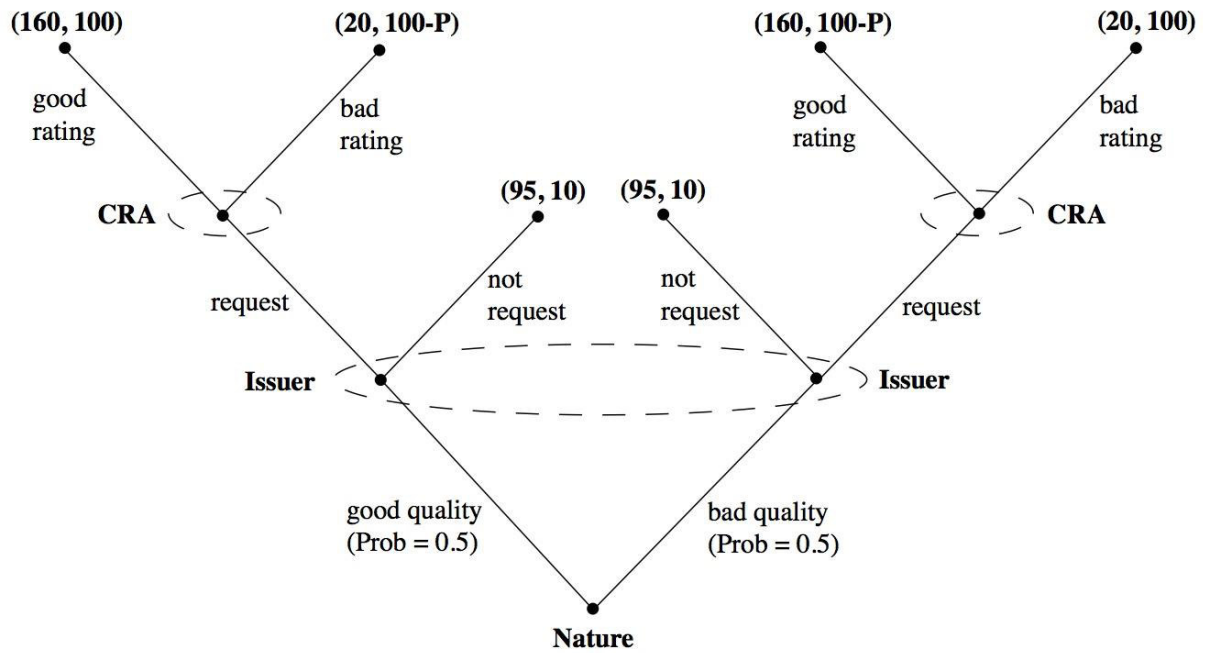


Figure 1: Decision tree

CRA, the payoffs are 95⁸ for the issuer and 10⁹ for the CRA. Figure 1 exhibits the sequence of decisions and the payoffs in a single period of the game.

Solving the game backwards, the issuer knows that, asked for a rating, the CRA would always report truthfully. The expected payoff is $0.5 * 160 + 0.5 * 20 = 90$ if the issuer requests a rating. This is lower than the payoff of 95 that the issuer gets if he does not request. Therefore, under the assumption of risk-neutrality we have a unique subgame perfect equilibrium solution, in which the CRA would always report truthfully and the issuer does not request a rating.¹⁰

The finite repetition of this game, for which we predict based on the backward induction argument that the issuer never requests a rating might bring to mind the “chain store paradox” of Selten (1978). The paradox is caused by the game-theoretical solution being unintuitive. Selten (1978) distinguishes between what he calls the “induction theory” and the “deterrence”, or more generally, the “cooperation theory”. While the induction theory assumes players to always stick to the rational, game theoretical solution in each repetition of the finite game, the cooperation theory allows players to strategically deviate from the game theoretical solution, as this might pay off to each player in the long run. In the finitely repeated version of our game, both issuer and CRA might have an incentive to follow a cooperative strategy. If we assume that the issuer requests and the CRA always provides a good report, the payoff to the issuer is 160, while the expected payoff to the CRA is $100 - P/2$ per period. Note that the latter is larger than 10 as long as the pecuniary injury is below 180. Thus, an issuer who requests a rating might

⁸ This amount is slightly higher than the expected payoff if the CRA always reported the true quality and thus reflects the saving of rating costs.

⁹ CRAs receive more than 90% of their income from issuer-fees (e.g., Darcy, 2009) and thus only about 10% from other activities.

¹⁰ Also a risk-averse issuer with $0.5u(160) + 0.5u(20) < u(90)$ would always choose not to request a rating since $0.5u(160) + 0.5u(20) < u(90) < u(95)$ by monotonicity of utility.

obtain an untruthful good rating, as long as the CRA has an incentive in being requested also in future periods.^{11,12}

2.2 Organization of the Experiment

Eighty-six subjects from different disciplines at the Georg-August-University of Göttingen participated in two treatments of the experiment, called the low-penalty and the high-penalty treatment. Producing an untruthful report implies relatively small pecuniary injuries in the low-penalty treatment with $P=10$, and relatively high pecuniary injuries in the high-penalty treatment with $P=50$. We conducted two sessions for each treatment and collected 21 independent observations in the low-penalty and 22 in the high-penalty treatment. Subjects were recruited via ORSEE (Greiner, 2004). The experiment was programmed and conducted via z-Tree (Fischbacher, 2007). The experiment sessions took place at the end of March and the beginning of April, 2015. Of the 86 participants, 36 percent were female and 45 percent were students of economics or business.

Issuers and CRAs interacted over 20 periods in a partner matching. An experiment session lasted around 75 minutes including the instructions and

¹¹ Cooperation could also be the (sequential equilibrium) outcome of a Bayesian reputation model (e.g., Kreps and Wilson, 1982; Kreps et al., 1982). Similarly to Frenkel (2015), assume that the CRA can be of two types, the truthful type always reporting the true quality and the cooperative type always reporting a good quality (with the exception of a potential “end play”). Let μ be the issuer’s belief that the CRA is of the cooperative type. The expected payoff to the requesting issuer is then given by $0.5 * 160 + 0.5 * (160\mu + 20(1-\mu))$, which is greater than the payoff of 95 in the case of no request if $\mu > 1/14$. Given an initial belief μ_0 , the issuer knows that the CRA is of the cooperative type ($\mu = 1$) if the CRA provided an untruthful good report or of the truthful type ($\mu = 0$) if the CRA provided a truthful bad rating; otherwise the issuer cannot update the belief. Until near the end of the interaction, as long as μ is above the threshold of $1/14$ requesting a report will be rational.

¹² A similar situation can be found in the centipede game (Rosenthal, 1981), where the second mover has an incentive to cooperate with the first mover since, once called upon to move, the second mover has reason to assume that the first mover is not following the (selfish) subgame perfect equilibrium strategy, but rather a strategy that is in the interest of both players.

payments. The instructions were read aloud in front of all subjects.¹³ We used a framing that was unloaded in that we did not talk about issuer and rating agency. The average amount earned (incl. a show-up fee of five Euros) by issuer-players is 16.88 Euros in the low-penalty and 16.44 Euros in the high-penalty treatment. The average amount earned by CRA-players is 12.48 Euros in the low-penalty and 11.39 Euros in the high-penalty treatment. Subjects were informed that they are not allowed to communicate, that questions are answered only privately, and that they will get paid directly after the experiment.

2.3 Behavioral Hypotheses

We derive behavioral hypotheses from the intuitive cooperation theory presented above. If a rating is requested, the CRA has, at least in early periods, an incentive to provide a good rating even if the true quality is bad since he has an interest in a rating being requested also in future periods (H2a). Anticipating this, the issuer has an incentive to request a rating, at least in the early periods (H1a). The CRA's interest in cooperation is stronger in the low-penalty case than in the high-penalty case. Thus the effects hypothesized in H1a and H2a should be stronger in the low-penalty case (H1b, H2b). The reasoning behind H1 and H2 should end in the final period(s); thus the anticipation of an end-game effect (Selten and Stoecker, 1986) (H3a, b).

H1: (a) We shall observe issuer-players to request ratings in both treatments.
(b) The relative frequency of requests will be smaller when the cost of untruthful reporting is higher.

¹³ The instructions can be found in the Appendix.

H2: (a) CRA-players, who are asked for a rating, show a tendency to provide a good rating even though the true quality is bad. (b) This tendency will be weaker, when the cost of untruthful reporting is higher.

H3: Report requests (a) and good ratings even though the true quality is bad (b) shall vanish toward the end of the interaction (end-game effect).

3. Experimental Results

Figure 2 shows the relative frequency with which issuers requested a rating both in the case of low and high costs of untruthful reporting. In each period, contrary to the equilibrium prediction, more than 50 percent of the issuers did request a rating. This holds true in both treatments. On average, the relative frequency of requested ratings amounts to 78.81 percent in the low-penalty and 70.68 percent in the high-penalty treatment. There is no statistically significant difference in the issuers' requesting behavior between the two treatments (U-Test, $p=0.2654$).¹⁴ We neither observe a pronounced downward trend nor a sharp decline in requests in the final period: 71.43 (54.55) percent of issuers request ratings in the last period in the low-penalty (high-penalty) treatment.¹⁵

Our results support H1a but they do not support H1b and H3a.

¹⁴ Using STATA for our statistical analysis, we denote the Mann-Whitney-U-Test as "U-test" and the Wilcoxon Signed Ranks Test as "WSR-test". We use two-sided versions of these tests.

¹⁵ The relative frequency of requests in periods 1 to 10 is not different from the one in periods 11 to 20 (WSR-test, $p=0.3056$ for the low-penalty and $p=0.6101$ for the high-penalty treatment). In both treatments, the relative frequency of requested ratings in the final period (period 20) of the game is not significantly different from any of the three previous periods, i.e. periods 17, 18, and 19 (McNemar tests, $p>0.1$).

When asked for a rating, we observe that, in good-quality states, CRAs provide truthful reports in 92.86 percent of all cases in the low-penalty treatment and in 97.90 percent of all cases in the high-penalty treatment. In bad-quality states, however, CRAs provide truthful reports only in 37.70 percent of all cases in the low-penalty and in 47.62 percent of all cases in the high-penalty treatment.

Considering our general hypothesis that issuers and CRA-players follow a cooperative strategy, the untruthful reporting of CRAs in bad-quality states (62.30 percent in the low-penalty and 52.38 percent in the high-penalty treatment) is of particular interest. Figure 3 shows the relative frequency of untruthful reports in each of the bad-quality states in the case of low and high costs of untruthful reporting.




Figure 2: Relative frequency of requested reports




Figure 3: Relative frequency of untruthful reports in bad-quality states

While, over all periods, there is no difference in untruthful reporting in bad-quality states between the two treatments (U-test, $p=0.1664$), there is a significant difference in the first bad-quality state of the game: in Period 2, the relative frequency of untruthful reports is significantly higher in the low-penalty case (Chi-squared test, $p=0.0200$). In both treatments, the relative frequency of untruthful reports in period 20 is significantly lower than in each of the three previous periods with bad-quality states, i.e. periods 13, 14, and 16 (McNemar tests, $p<0.01$).¹⁶

Our results support H2a and H3b but they do not support H2b.

Over the course of the interaction, the CRA's decision to cooperate with the issuer and to provide untruthful ratings in bad-quality states might have an impact on the decisions of issuers to request ratings. We test this hypothesis by means of a regression, of which the results are presented in Table 2. In the specified regression model, we investigate whether the issuer's decision to request a rating is influenced by the last experience with the CRA when the quality-state was good and the one when the quality-state was bad. The results are straightforward: reporting untruthfully in a good-quality state decreases and reporting untruthfully in a bad-quality state increases the probability of a rating being requested again. The latter implies that, given the opportunity to provide a rating, the CRA might encourage the issuer to request ratings in the future by providing an untruthful report in bad-quality states.¹⁷

¹⁶ We consider only groups requesting a rating in both periods considered.

¹⁷ Both issuer and CRA earned significantly higher payoffs than the respective equilibrium payoff. While issuers earned an average payoff per period of 109.04 (104.68), CRAs earned an average payoff per period of 77.86 (63.27) in the low-penalty (high-penalty) treatment. Each of the CRA-players earned more than the equilibrium payoff of 10. Comparing the number of issuers who earned payoffs that are higher than the equilibrium payoff of 95 (15 in the low-penalty and 18 in the high-penalty treatment) to the number of issuers who earned payoffs that are lower or equal to 95 (6 in the low-penalty and 4 in the high-penalty treatment), we find that cooperation significantly pays off to issuers as well (two-sided Binomial tests, $p=0.0784$ for the low-penalty and $p=0.0043$ for the high-penalty treatment).

Table 2: The impact of previous experience on the issuers' requesting behavior

		Dependent variable: request	
		(1) Low Penalty	(2) High Penalty
Last rating received in a good-quality state	Bad rating	-0.7109** (0.3291)	-1.3054** (0.2396)
	Good rating	Ref.	Ref.
Last rating received in a bad-quality state	Bad rating	Ref.	Ref.
	Good rating	1.0040*** (0.2164)	1.0334*** (0.2396)
Constant		0.5196*** (0.2086)	0.4492* (0.2364)
# of subjects		21	22
# of observations		366	361

Random-effect probit model; standard-errors are in parenthesis
 *** Significance at 1%, ** 5%, * 10%

4. Conclusion

In this study, we introduced a simple game-theoretical model that captures the main aspects of the interaction between an issuer and a CRA under relatively strict assumptions. We parameterized the model such that, in the subgame-perfect equilibrium, CRAs should neither produce optimistically nor pessimistically biased ratings and, knowing this, issuers should not request ratings. Conducting laboratory experiments, however, we find that in about three quarters of all rounds ratings are requested. This is reciprocated with more than 50 percent of inflated ratings, even though CRAs are penalized for

untruthful reports. We consider a low and a high level of pecuniary injuries caused by untruthful reports, where even the high cost level turns out not important enough to make cooperation uninteresting.

Our findings could have implications for the policy design in financial market regulation. While there have been several attempts to regulate the credit rating industry like, among others, the Credit Rating Agency Reform Act of 2006, the Dodd-Frank Wall Street Reform and Consumer Protection Act, or Andrew Cuomo's agreement, there is very little research on the effectiveness of those reforms.¹⁸ Our study may be seen as an attempt to evaluate the impact of Cuomo's agreement in a simplified monopolistic world, where the agreement is strictly implemented (i.e., no possibilities for communicating and reaching hidden agreements, direct publication of any rating that has been requested, financial penalties or reputational costs for inaccurate ratings). It shows to be little effective, since it is most likely that issuers and CRAs prefer to cooperate rather than to stick to the rules. This result should also hold in an oligopolistic rating market, where an issuer receiving an unsatisfactory rating today has the opportunity to solicit future ratings from different CRAs. In that case, a CRA would have an even higher interest in establishing a long-term cooperation with an issuer than in a monopolistic rating market.

¹⁸ The recent study of Dimitrov et al. (2015) is one of the few that evaluates a regulatory reform on CRAs. We find no studies that explicitly focus on the evaluation of the Credit Rating Agency Reform Act of 2006.

References

Ashcraft, A. B., Goldsmith-Pinkham, P. and Vickery, J. I. (2010). "MBS Ratings and the Mortgage Credit Boom". Federal Reserve Bank of New York Staff Report 449.

Becker, B. and Milbourn, T. (2011). "How did Increased Competition Affect Credit Ratings?". *Journal of Financial Economics*, Vol. 101, Issue 3, pp. 493-514.

Bolton, P., Freixas, X. and Shapiro, J. (2012). "The Credit Ratings Game". *The Journal of Finance*, Vol. 67, Issue 1, pp. 85-111.

Bongaerts, D. (2014). "Alternatives for Issuer-Paid Credit Rating Agencies". ECB Working Paper Series, No. 1703.

Camanho, N., Deb, P. and Liu, Z. (2012). "Credit Rating and Competition". 22nd Australasian Finance and Banking Conference, AFA 2011 Denver Meetings Paper.

Darcy, D. (2009). "Credit Rating Agencies and the Credit Crisis: How the "Issuer Pays" Conflict Contributed and What Regulators Might Do About It". *Columbia Business Law Review*, Vol. 2009, No. 2:605, pp. 605-668.

Deb, P. and Murphy, G. (2009). "Credit Rating Agencies: An Alternative Model". Working paper.

Dimitrov, V., Palia, D. and Tang, L. (2015). "Impact of the Dodd-Frank Act on Credit Ratings". *Journal of Financial Economics*, Vol. 115, Issue 3, pp. 505-520.

Financial Crisis Inquiry Commission (2011). "The Financial Crisis Inquiry Report". Available at: <https://www.gpo.gov/fdsys/pkg/GPO-FCIC/pdf/GPO-FCIC.pdf>. Retrieved on January 27, 2016.

Fischbacher, U. (2007). "z-Tree: Zurich Toolbox for Ready-Made Economic Experiments". *Experimental Economics*, Vol. 10, Issue 2, pp. 171-178.

Frenkel, S. (2015). "Repeated Interaction and Rating Inflation: A Model of Double Reputation". *American Economic Journal: Microeconomics*, Vol. 7, No. 1, pp. 250-280.

Greiner, B. (2004). "An Online Recruitment System for Economic Experiments". *Forschung und Wissenschaftliches Rechnen 2003*, GWDG Bericht, Vol. 63, pp. 79-93.

Griffin, J. M. and Tang, D. Y. (2012). "Did Subjectivity Play a Role in CDO Credit Ratings?" *The Journal of Finance*, Vol. 67, Issue 4, pp. 1293-1328.

Jiang, J. (X.), Stanford, M. H. and Xie, Y. (2012). "Does it Matter Who Pays for Bond Ratings? Historical Evidence". *Journal of Financial Economics*, Vol. 105, Issue 3, pp. 607-621.

Kreps, D. M. and Wilson, R. (1982). "Reputation and Imperfect Information". *Journal of Economic Theory*, Vol. 27, Issue 2, pp. 253-279.

Kreps, D. M., Milgrom, P., Roberts, J. and Wilson, R. (1982). "Rational Cooperation in the Finitely Repeated Prisoners' Dilemma". *Journal of Economic Theory*, Vol. 27, Issue 2, pp. 245-252.

New York State Office of the Attorney General (2008). "Attorney General Cuomo Announces Landmark Reform Agreements With The Nation's Three Principal Credit Rating Agencies". Available at: <http://www.ag.ny.gov/press-release/attorney-general-cuomo-announces-landmark-reform-agreements-nations-three-principal>. Retrieved on February 11, 2016.

Rosenthal, R. (1981). "Games of Perfect Information, Predatory Pricing and the Chain-Store Paradox". *Journal of Economic Theory*, Vol. 25, Issue 1, pp. 92-100.

Securities and Exchange Commission (SEC) (2008). "Summary Report of Issues Identified in the Commission Staff's Examinations of Select Credit Rating Agencies". Available at: <https://www.sec.gov/news/studies/2008/craexamination070808.pdf>. Retrieved on January 26, 2016.

Selten, R. (1978). "The Chain Store Paradox". *Theory and Decision*, Vol. 9, Issue 2, pp. 127-159.

Selten, R. and Stoecker, R. (1986). "End Behavior in Sequences of Finite Prisoner's Dilemma Supergames A Learning Theory Approach". *Journal of Economic Behavior & Organization*, Vol. 7, Issue 1, pp. 47-70.

Skreta, V. and Veldkamp, L. (2009). "Ratings Shopping and Asset Complexity: A Theory of Ratings Inflation". *Journal of Monetary Economics*, Vol. 56, Issue 5, pp. 678-695.

Strobl, G. and Xia, H. (2011). "The Issuer-Pays Rating Model and Ratings Inflation: Evidence from Corporate Credit Ratings". Working paper. Available at SSRN: <http://ssrn.com/abstract=2002186> or <http://dx.doi.org/10.2139/ssrn.2002186>

White, L. J. (2010). "Markets: The Credit Rating Agencies". *The Journal of Economic Perspectives*, Vol. 24, No. 2, pp. 211-226.

Appendix – A: Instructions (low penalty, translated from German)

Welcome to this Experiment!

From now on, and until the end of the experiment, communicating with other participants is prohibited. If you have any question(s), please raise your hand. We will come to you directly and answer your question(s) individually.

The game plays over **20 rounds**. There are two types of players: **players A** and **players B**. You are randomly assigned at the beginning of the game to one of these two roles (player A or player B). You keep your role over the entire game. Every player A is randomly matched to a player B at the beginning of the game and will interact with the same person over all 20 rounds. You will not be informed about the identity of the person, you were matched with, at any point in time.

The Decisions

Player A owns a product of unknown quality. The quality of this product can be (each with a probability of **50%**) either **good** or **bad**. The quality of the product is randomly determined at the beginning of each round. Neither player A nor player B are informed about the quality of the product.

A product of unknown quality has a value of **95 ECU (Experimental Currency Unit)**. In case player A does not request a report on the value of the product, the round ends immediately and player A receives a payment of 95 ECU.

Player A can receive a different payment if she requests a report about the quality of the product from player B. **Player A** has to decide in every round whether or not to request a report from player B.

In case player A requests a report from player B, **player B** obtains a **private information** that reveals the quality of the product. **Independently** of the private information player B then decides to send one of the following reports:

- The product is of **good** quality
- The product is of **bad** quality

If player B decides to send the report “the product is of **good** quality”, player A will receive a payment of **160 ECU** for her product. If player B decides to send the report “the product is of **bad** quality”, player A will receive a value of **20 ECU** for her product. The round then ends.

The payment of **player B** in a round is computed as follows:

- In case player A did not request a report, then player B receives **10 ECU**.
- In case player A requested a report from player B, and this report corresponded to the private information received, then player B receives **100 ECU**.
- In case player A requested a report from player B, and this report did not correspond to the private information received, then player B receives **90 ECU**.

The **total payoff** of each player is **the sum of the payments in each round**. This payoff in ECU will be converted to Euros at the following rate: 10 ECU = 5 euro cents. At the end of the 20 rounds, you will receive the

amount earned in the experiment, to which will be added an additional fixed payment of €5.

Information Available during the Experiment

In each round an overview table showing the outcomes of all previous rounds will be displayed on your screen. The table includes for each round the following information: the decision of player A, the report of player B, the randomly determined quality of the product, the payoff of player A, the payoff of player B.

Please go to the computer corresponding to your participant number. After clicking the button “continue”, some questions will be displayed to test your understanding of the instructions. The game starts only once all the questions have been answered correctly by all participants. Please raise your hand, should you have any questions. We will then come directly to you and answer your question individually.