ecgi

A Strict Liability Regime for Rating Agencies

Law Working Paper N° 245/2014
March 2014

Alessio M. Pacces
Erasmus University Rotterdam and ECGI

Alessandro Romano
Erasmus University Rotterdam and LUISS

ECGI Working Paper Series in Law

A Strict Liability Regime for Rating Agencies

Working Paper N°. 245/2014
March 2014

Alessio M. Pacces
Alessandro Romano

This paper benefited from the feedback of participants in the 2013 Meeting of the Canadian Association in Law and Economics in Toronto, in the 15th Joint Seminar of the European Association of Law and Economics and The Geneva Association in Girona, in the 4th Annual Conference of the Spanish Association of Law and Economics in Granada, and in the 11th Annual Meeting of the German Association of Law and Economics in Bolzano. We are particularly indebted to Gaia Barone, Alice Guerra, Claire Hill, Giovanni Immordino, Vyacheslav Mikhed, Marcello Puca, Pietro Reichlin, Angela Troisi, and Roland Vaubel for their valuable comments.

Abstract

This paper argues that a mitigated strict liability regime can incentivize Credit Rating Agencies (CRAs) to produce ratings as accurate as the available forecasting technology allows. A damage cap based on objective factors is introduced in order to avoid crushing liability. Moreover, CRAs are allowed to choose how much to commit to their predictions. CRAs may opt out of liability even entirely, unless their ratings are relevant for regulation. Finally, corrections in the relevant timeframe for the imposition of liability are introduced in order to protect CRAs from systemic risk.

Keywords: law & economics; financial regulation; rating inflation; probability of default; crushing liability; imperfect foresight; systemic risk; structured finance

JEL Classifications: G01; G24; K13; K22

Alessio M. Pacces*
Professor of Law and Finance
Erasmus University Rotterdam, School of Law
Room L05-119, P.O. Box 1738
3000 DR Rotterdam, The Netherlands
phone: +31-(0)10-4081613, fax: +31-(0)10-4089191
e-mail: pacces@law.eur.nl

Alessandro Romano
Ph.D. Student in Law and Economics
Erasmus University Rotterdam, School of Law
Burgemeester Oudlaan 50
3062 Rotterdam, The Netherlands
e-mail: aromano@live.it

*Corresponding Author

A Strict Liability Regime for Rating Agencies

Alessio M. Pacces[♣], Alessandro Romano[♦]

(Version: 6 March 2014)

Abstract

This paper argues that a mitigated strict liability regime can incentivize Credit Rating Agencies (CRAs) to produce ratings as accurate as the available forecasting technology allows. A damage cap based on objective factors is introduced in order to avoid crushing liability. Moreover, CRAs are allowed to choose how much to commit to their predictions. CRAs may opt out of liability even entirely, unless their ratings are relevant for regulation. Finally, corrections in the relevant timeframe for the imposition of liability are introduced in order to protect CRAs from systemic risk.

Keywords: law & economics; financial regulation; rating inflation; probability of default; crushing liability; imperfect foresight; systemic risk; structured finance.

JEL Classification: G01; G24; K13; K22.

1. Introduction

The behavior of credit rating agencies (henceforth CRAs) has been under careful scrutiny in the past decade, particularly in the aftermath of the global financial crisis. It has been argued that the incentives of CRAs are adversely affected by an inherent conflict of interest determined by the "issuer-pays model" (Krugman 2010; Pagano and Volpin 2010) and by the licensing power that financial regulations relying on ratings implicitly grant to CRAs (Partnoy 2010; Calomiris 2009a; Opp, Opp and Harris 2013). In this perspective, ratings

[♣] Erasmus University Rotterdam (School of Law) and European Corporate Governance Institute - pacces@law.eur.nl

[♦] Erasmus University Rotterdam (School of Law) and LUISS Guido Carli (School of Economics) - aromano@live.it

This paper benefited from the feedback of participants in the 2013 Meeting of the Canadian Association in Law and Economics in Toronto, in the 15th Joint Seminar of the European Association of Law and Economics and The Geneva Association in Girona, in the 4th Annual Conference of the Spanish Association of Law and Economics in Granada, and in the 11th Annual Meeting of the German Association of Law and Economics in Bolzano. We are particularly indebted to Gaia Barone, Alice Guerra, Claire Hill, Giovanni Immordino, Vyacheslav Mikhed, Marcello Puca, Pietro Reichlin, Angela Troisi, and Roland Vaubel for their valuable comments.

are inflated either because issuers collude with CRAs in fooling investors or because, all else being equal, investors demand assets with higher ratings in order to enjoy regulatory benefits. Inflated ratings, the argument runs, support asset bubbles, which are in turn a major determinant of financial crises (Calomiris 2009b). Although we are agnostic about the actual contribution of ratings to the global financial crisis, we follow the mainstream literature assuming that accurate ratings are valuable for the society (Bolton, Freixas Shapiro 2012), whereas inflated ratings may reduce welfare, particularly when ratings have regulatory relevance.

We acknowledge that ratings are ultimately predictions and thus they can be as accurate as our ability to forecast the future can be. This observation has important consequences on how, in this paper, we argue that the incentives of CRAs should be policed. CRAs should be in principle allowed to choose how much to commit to the accuracy of their prediction, if to commit at all. That being said, we define inaccurate a rating whose implied predictions is not borne out by the actual unfolding of events. To simplify, a rating with a certain letter grade (for example Double-A+) is inaccurate if the frequency of default of firms or bonds with that letter grade is higher or lower respectively than the maximum (for example 0.0006) and the minimum (for example 0.0002) probability of default associated with the letter grade.[1] When the frequency of defaults turns out to be higher than the upper bound on the predicted probability of default, we speak of rating inflation.[2] A rating is accurate when the defaults actually observed for a given class of rating fall within the range of probabilities and other measurable items (for instance, loss given default) implied by the CRA issuing a certain letter grade.[3]

In this paper, we argue that the accuracy of ratings can be improved via regulatory intervention, particularly by introducing a special liability rule for CRAs. This approach has been little explored by the literature. Apparently, a more straightforward solution to the problem of rating inflation could be based on eliminating its determinants by regulation. In this vein, all references to credit ratings could be scrapped from financial regulation in

[1] The example is taken from Fitch's historical (annualized) default experience. See Coval, Jurek and Stafford (2009b).
[2] For a formal definition of rating inflation, see section 4.
[3] In this paper we do consider the opposite reason of inaccuracy, namely rating deflation. However, for the reasons discussed in section 4, addressing this problem is not so interesting for policymaking. Under the status quo, where CRAs hardly face any liability, we argue that CRAs always have incentives to inflate ratings. Introducing liability may induce CRAs to systematically underrate financial assets. However, at some point this strategy would make ratings uninteresting for issuers and investors. Section 5.2 explicitly discusses why inducing CRAs to be moderately conservative with ratings is desirable, particularly in the case of structured finance products.

order to eliminate the regulatory benefits from high ratings (Flannery, Houston and Partnoy 2010). This is, incidentally, the approach chosen by the U.S. legislation with the Dodd-Frank Act of 2010.[4] Likewise, it could be argued along with a number of commentators (see Coffee 2006; Mathis, McAndrews and Rochet 2009) that the issuer-pays model of CRAs remuneration is simply to be prohibited in order to eradicate the conflict of interests.

As straightforward as they may sound, these radical proposals of regulatory intervention are too farfetched. Rating agencies play a crucial role in helping to overcome information asymmetries not only between issuers and investors, but also between the latter and financial regulators. In the absence of viable alternatives to assess creditworthiness and credit risk, it is at least doubtful that financial regulation could just do without ratings (Coffee 2010). Similarly, the public good nature of ratings – the use of ratings does not diminish their availability to others; and investors who do not pay for ratings can hardly be excluded from their use – might frustrate the attempt to introduce a workable alternative to the issuer-pays model (Partnoy 1999). More importantly, moving away from the issuer-pays model would not solve the problem so long as regulatory benefits are present. Because at least some regulated investors demand high ratings irrespective of their informativeness, switching to an investor-pays model is unlikely to stop rating inflation.

Abandoning the realm of radical reforms, even more modest changes of the status quo proposed so far seem to suffer from serious drawbacks. For example, let us consider two of the most popular incremental reforms in the policy debate. One proposal is to increase competition between CRAs (Partnoy 2010). The other is to increase the transparency of their ratings (Pagano and Volpin 2010). Both reforms aim at reducing the ability of CRAs to collude with issuers or investors to generate inflated ratings. However, competition between CRAs is set to make matters worse because of the practice of so-called "rating shopping". Because issuers can solicit[5] as many ratings as they wish but pay for rating only if they request publication, more competition between CRAs may actually result in more rating inflation (Becker and Milbourn 2011). To be sure, rating shopping could be

[4] Section 939A of the Wall Street Reform and Consumer Protection Act of 2010 (the Dodd-Frank Act) requires each Federal agency to remove references to credit rating. The implementation of this provision has proven difficult, although major agencies like the Federal Reserve Board and the Securities and Exchange Commission have ultimately found ways to issue the necessary regulations. The approach in the EU has been different. While EU legislation also aims at reducing over-reliance on ratings (see the Capital Requirements Directive IV and the recent Regulation 462/2013 and Directive 2013/14/EU on credit ratings), it explicitly acknowledges that financial regulation cannot simply do away with ratings in the absence of viable alternatives.

[5] This paper does not deal with unsolicited ratings, which typically concern sovereign issuers.

prohibited, for example by requiring issuers to pay for ratings in advance[6] and CRAs to disclose also unfavorable ratings. This solution may not solve the problem of implicit rating shopping, though, as issuers could learn the CRAs' assessment informally before entering into a contract with them (Sangiorgi, Sokobin and Spatt 2009; Pagano and Volpin 2010). At the same time, forcing issuers to pay for ratings without knowing their contents may generate moral hazard. If CRAs can save on their costs after having secured an income independent of their assessment, eventually this would lead to the collapse of the market for ratings (Bolton, Freixas and Shapiro 2012).

The economic literature on CRAs has been so far unable to identify a workable policy, whether radical or incremental, that could ameliorate the incentive problems leading to rating inflation. However, the problem is in principle a simple one to solve: CRAs should earn market profits from producing accurate ratings but be punished if they produce inflated ratings, at least inasmuch as this behavior results in negative externalities to society. Since Coase (1960) and Calabresi (1970), law and economics identifies in the legal liability one of the instruments for policing incentives to produce negative externalities. In the presence of negative externalities, liability can improve welfare if transaction costs are sufficiently high to prevent market forces from coping with the problem.

Our paper follows this law and economics tradition. In the context of CRAs, transaction costs are high when reputational concerns are insufficient to stop the production of inflated ratings. However, because this is a classic commitment problem, we argue that it can be improved by appropriate enforceable contracts (Cooter and Ulen 2011), including liability for ratings that turn out to be inaccurate. In this case CRAs should be able to *choose* how much exposure to liability is necessary to commit to levels of accuracy that investors (and thus issuers) find acceptable to sustain a market for ratings. The situation is different when reputation is not just insufficient to commit CRAs to a level of accuracy of their choice, but is displaced altogether by the ability of CRAs to support regulatory arbitrage, for instance because investing in a Triple-A asset of whatever creditworthiness brings regulatory benefits. In this situation, it is impossible to put the Coase Theorem back to work. On the one hand, CRAs are unambiguously better off opting out of any liability. On the other hand, those who suffer from inflated ratings (for example unregulated investors

[6] This is the essence of the so-called Cuomo Plan, named after the New York State Attorney General who proposed this approach. As noted by Bolton, Freixas and Shapiro (2012), this approach does not eliminate rating shopping in the absence of an explicit obligation to disclose also unfavorable ratings.

fooled by high ratings; or taxpayers bearing the cost of bailouts) can hardly negotiate with CRAs a commitment to accurate ratings even if that would improve welfare. In this case, the market is unable to correct the negative externalities problem. Hence, regulation should set a minimum degree of exposure to liability as a condition for ratings to enjoy regulatory relevance.

In order to improve the incentives of rating agencies we advocate the introduction of a simple and legally workable strict liability rule: CRAs should be liable to pay damages whenever a bond or a company they rate defaults. This is different from the approach taken by regulation on both sides of the Atlantic in the aftermath of the global financial crisis. While in the US and, more recently, in the EU, CRAs have been subject to liability based on negligence (if not gross negligence or even intent),[7] we argue that CRAs should face strict liability with three strong limitations. First, damage compensation should be capped at a multiplier of the CRA's income. Second, liability should operate with a timeframe apt to shield CRAs from systemic risk. Third, at least in the absence of regulatory benefits, CRAs should be able to decide how much to commit to their ratings by choosing a certain degree of liability exposure.

We set these limitations in order to avoid crushing liability (Shavell 1980). Crushing liability deters a socially valuable activity, like the production of accurate ratings, by imposing on the actor subject to it a liability in excess to the harm that it causes to the society. We argue that strict liability would be crushing for CRAs if they were liable for more than their revenues from selling ratings that are as accurate as possible, given the limits of the existing forecasting models as reflected by the chosen level of commitment. Likewise, crushing liability would stem from correlated defaults requiring CRAs to pay damages, however capped. For simplicity, we call these correlated defaults systemic risk. Systemic risk cannot be insured and, because CRAs are effectively silent about systemic risk, they should not be responsible for it[8].

[7] In the U.S., the exemption of CRAs from liability as experts pursuant to Section 11 of the Securities Act of 1933 was removed in 2010 (see Dodd-Frank Act § 939G). As a result, CRAs are currently subject to liability under a due diligence standard provided that they are named as experts in the prospectus, which they can and do refuse (see Coffee 2010). On this side of the Atlantic, a EU-wide liability of CRAs was only introduced in 2013. "Where a credit rating agency has committed, intentionally or with gross negligence, any of the infringements listed in Annex III having an impact on a credit rating, an investor or issuer may claim damages from that credit rating agency for damage caused to it due to that infringement" (art. 35a, 1, Reg. (EC) no. 1060/2009 as amended by art. 1, (22), Reg. (EU) no. 462/2013).

[8] As noted by Coval, Jurek and Stafford (2009b: 23), credit ratings "are silent regarding the state of the world in which default is likely to happen." Therefore, ratings are uninformative about systemic risk.

We show with a reduced-form model that strict liability leads CRAs to produce more accurate ratings under the three limitations sketched out above. To begin with, the damages are capped based on the income from rating divided by the highest probability of default associated with the letter grade of the defaulted asset. This condition is sufficient to disallow profits from rating inflation without discouraging ratings altogether. More precisely, CRAs facing this strict liability make no loss conditional on the absence of rating inflation as revealed by the difference between the predicted default rate and the actual frequency of defaults.

Moreover, a correction is introduced to protect CRAs from defaults depending on systemic risk. We suggest two different approaches for corporate bonds and for structured finance products, because they have a very different exposure to systemic risk. We note that corporate defaults tend to be strongly correlated only in the medium to long term. Therefore, as far as corporate bonds are concerned, we argue that liability should operate only for a limited period after the production or the confirmation of a rating. Although this is sufficient for corporate bonds, the defaults of structured finance products tend to be correlated also in the short term, particularly in a financial crisis (Coval, Jurek and Stafford 2009a,b). Because in this situation strict liability may discourage CRAs from rating structured finance altogether, we propose another solution to cope with systemic risk. Whenever extraordinary default rates are arguably dependent on systemic risk, liability should be conditional on inaccuracy being confirmed by the law of large numbers. When a public authority announces a financial crisis status, liability would be imposed on CRAs only if the frequency of observed defaults departed from the predictions made by CRAs over a sufficiently large number of cases and a sufficiently large time span, thus protecting CRAs from violent short-term fluctuations in the default rates. While limiting the extent to which strict liability over-deters ratings, particularly of structured finance products, this solution is countercyclical as it rewards the CRAs that were more conservative in their assessments during the upswing phase of an asset bubble.

Finally, CRAs are allowed to decide how much to commit to a certain rating, that is to the probabilities of default and the other estimates associated with each letter grade, by choosing the degree of exposure to liability. This condition allows liability to reflect the uncertainty of the forecasting models available to CRAs. The limited ability to foresee the future, along with the unobservability of several variables affecting the performance of the market for rating, is the reason why we ultimately advocate a contractual approach to

CRAs liability. This approach, however, leaves us with a problem. In the presence of regulatory benefits, CRAs may choose an inefficiently low level of commitment and profit from providing regulated investors with artificially high ratings (Opp, Opp and Harris 2013). To address this issue, regulation should require that CRAs face a minimum degree of liability exposure for their rating to enable regulatory benefits. This solution would still allow CRAs to choose their commitment levels with investors, but only in the absence of the negative externalities created by inflated ratings with a regulatory value.

The reminder of the paper is as follows. In section 2 we discuss the literature related to the credit rating problem and show how our paper contributes to it. In section 3 we introduce our strict liability rule with a simple numerical example. In section 4 we illustrate in two phases, with a reduced-form model, how capped strict liability prevents CRAs from inflating ratings. First, we present this result under simplifying assumptions. Second, we show that the choice of liability exposure by CRAs can lead to the same result under the more realistic assumptions of imperfect foresight and unobservability of reputation effects and other key variables. In section 5 we discuss the problem of systemic risk and how the strict liability rule should be corrected to cope with this problem in the two cases of corporate bonds and structured finance products. In section 6 we illustrate the key advantages of our proposal in disciplining CRAs' behavior. We briefly conclude with section 7.

2. Related Literature

Our paper contributes to the literature on Credit Rating Agencies by identifying a workable legal policy which can counter the CRAs' incentives to produce inflated ratings. In the literature on CRAs, the existence of rating inflation is rarely disputed.[9] However, the causes underlying rating inflation are not settled and there are different theories in this regard.

According to a first strand of literature the fundamental reason why CRAs tend to inflate their ratings is that they are paid by the same issuers that they rate. In this vein, the problem of rating inflation would be solved if one could simply make investors pay for

[9] One important exception is Gorton and Ordonez (2014) citing inter alia the study by Park (2011). This study, however, does not deny that the triple-A subprime-related securities turned out to be riskier than implied by their initial rating. Rather, their point is that few of these securities actually defaulted and that the losses stemming from such defaults were quantitatively small (too small to justify a global financial crisis).

ratings, which is complicated by information leakage and the related free riding problem (Pagano and Volpin 2010). However, even if it were possible to do away with the issuer-pays model, the case for legal intervention would not be straightforward. In a well-functioning market, reputational sanctions and competitive pressure could prevent opportunistic behavior by CRAs, regardless of the paying scheme adopted.

Many theoretical models have been developed to demonstrate how rating inflation emerges under different assumptions, thus suggesting the existence of market failures. Bolton, Freixas and Shapiro (2012) show that rating inflation can be driven by investors' naivety and by the freedom granted to issuers to purchase the rating that they prefer, which allows for rating shopping. Because the marginal investors may be unsophisticated and thus unable to identify and punish inaccurate ratings, CRAs will face lower reputational sanctions from inaccuracy while profiting from selling inflated ratings to issuers. Skreta and Veldkamp (2009) emphasize that rating inflation might emerge also in the presence of truth-telling CRAs if there is sufficient heterogeneity in the predictions of their models. A similar point is made by Sangiorgi, Sokobin and Spatt (2009). They argue that heterogeneity in CRAs' predictions results in rating inflation even if explicit rating shopping is forbidden. The reason is that rating shopping can always occur implicitly. Because the methodologies of rating agencies are transparent to a certain extent, the issuer can select the CRA that uses model assumptions allowing for the highest possible rating.

Opp, Opp and Harris (2013) take a different approach and show that rating inflation can depend exclusively on the regulatory function assigned to the ratings. Because ratings are embedded in financial regulation worldwide, regulated investors benefit from investing in highly rated securities even if the ratings are inaccurate. This strategy, for example, may lower the regulatory capital requirements for banks; may protect institutional investors from the threat of liability; and so forth (Partnoy 2010). The underlying assumption is that the value of these regulatory benefits passed on to CRAs via the issuers' fees exceeds the reputational sanction stemming from inflated ratings. More specifically, if regulatory benefits of high ratings are above a certain threshold, a rating agency "finds it profitable to stop acquiring any information and merely facilitates regulatory arbitrage through rating inflation" (Opp, Opp and Harris 2013, p. 47). The implications of this approach are twofold. On the one hand, it is not necessary to assume investors naivety to explain inflated ratings. On the other hand, to the extent that inflated ratings depend on a demand

by regulated investors, having investors rather than issuer pay for them cannot possibly ameliorate the problem (Calomiris 2009a).

Although due to identification problems rating inflation is hard to show empirically, there is some empirical evidence suggesting its presence as well as its dependence on several market failures. Using a panel dataset covering from 1999 to 2009, Xia and Strobl (2012) find that the issuer-pays practice leads to higher ratings than the investor-pays practice. Baklyar and Galil (2011) gather empirical evidence on the Israeli corporate credit rating market and show that one agency (Midroog) systematically inflated ratings, whereas another (S&P-Maalot) inflated its ratings only as a response to rating shopping. Becker and Milbourn (2011) hint at rating inflation only indirectly. Their study reveals that the entry of Fitch in the market for ratings worsened the quality of ratings. This finding suggests that the adverse effects of rating shopping on rating inflation outweigh the benefits of increased competition.

The lesson to be learnt from the theoretical and the empirical literature is that a combination of market failures and regulatory distortions probably exists. Ratings tend to be inflated because there are naïve investors, which make reputation a weak constraint on rating shopping, *and* because there are regulatory benefits, which allow CRAs to cater to the investors' demand for artificially high ratings. Moreover, there is no easy way in which the market or regulation can overcome these problems. If the marginal investors are naïve the market cannot easily self-correct. Put differently, because transaction costs prevent efficient contracts on the provision of ratings from being written, the Coase Theorem breaks down. Regulation could paternalistically protect naïve investors by prohibiting the issuer-pays model, rating shopping, or even both of them. However, this approach would hardly be effective. On the one hand, in the absence of a regulator or a court who can screen rating quality, a market for ratings deprived of its typical features may collapse because of free riding (Pagano and Volpin 2010) or moral hazard (Bolton, Freixas and Shapiro 2012). On the other hand, so long as financial regulation lacks viable alternatives to ratings for assessing credit risk, ratings will still be inflated despite any prohibition of issuer-pays and/or rating shopping when the regulatory benefits from inflation are high enough.

We maintain, along with the mainstream literature, that ratings are valuable for the society because they reduce asymmetric information in finance (White 2010). However, this is conditional on ratings being above a certain accuracy threshold, which for simplicity

we assume to be exogenously determined by the existing forecasting technology. Based on the findings of the existing literature, we see two reasons why CRAs may produce inaccurate ratings. One is a commitment problem (Cooter and Ulen 2011). Ratings are inflated because investors at the margin cannot recognize and punish inaccurate ratings (or cannot reward only accurate ratings), which prevents CRAs from committing to accurate predictions. The other reason is the presence of negative externalities in financial markets (Heremans and Pacces 2012). Because financial regulation currently relies on ratings to cope with such externalities, inaccuracy of ratings adversely affects not only the investors purchasing the rated assets for regulatory benefits, but also their counterparties as well as the taxpayers who bear the costs of bailing out regulated investors. We further assume that neither unsophisticated investors nor financial regulators can second-guess the quality of ratings, as we do not see how CRAs could produce anything valuable otherwise. It follows that there is a case for a different kind of legal intervention than proposed so far. Rather than tampering with how the market for ratings works or scrapping the distortions stemming from financial regulation, we propose to subject CRAs to legal liability while keeping the rest of the status quo.

Unfortunately, precisely because it is difficult for a third party to second-guess ratings, it has been virtually impossible to prove in courts the negligent behavior of rating agencies and the portion of losses suffered by investors that is attributable to their conduct. Therefore, as shown by the legal literature (Deats 2010), CRAs have been de facto immune from liability claims. Moreover, particularly in the U.S., the rating agencies have often been able to escape liability by invoking the protection of the First Amendment available to journalists, whose liability is subject to an actual malice standard. Despite the efforts of legislators on both side of the Atlantic, this situation is not likely to change in the near future.[10] Acknowledging the difficulty to police the incentives of CRAs through a negligence standard for tort liability, the law and economics literature has suggested imposing a punishment on CRAs that produce inaccurate ratings by paying them with the debt they rate (Listokin 2010). We follow a similar logic, although we are more optimistic about the possibility to implement this logic through a strict liability regime. Our approach

[10] In the U.S., CRAs face now more difficulties to invoke the protection of the First Amendment. However, it has been practically impossible to activate the CRAs' liability as experts under Section 11 of the Securities Act of 1933 introduced by the Dodd-Frank Act (see Coffee 2010). With regard to Europe it is doubtful that the gross negligence standard that was introduced by art. 35a, Reg. (EC) no. 1060/2009 as amended by art. 1, (22), Reg. (EU) no. 462/2013 will change the status quo. In fact, it will be very hard to prove in courts (i) the grossly negligent behavior, (ii) the causation, and (iii), the portion of the losses suffered by investors that is attributable to the CRA's conduct. Haar (2013) provides a comprehensive comparative discussion of the recent legal developments concerning the civil liability of CRAs.

has the important advantage to allow corrections for systemic risk, which are obviously not available for debt. We build on one of the policy recommendations by Bolton, Freixas and Shapiro (2012) to fix the weakness of CRAs' reputational constraint, namely the enhancement of legal liability. However, differently from them as well as from the rest of the literature, we design a complete liability regime supporting the production of accurate ratings (as accurate as allowed by the available forecasting technology) without undermining the existence and the functioning of a market for ratings.

The function of CRAs is to provide investors with certifications of the quality of financial assets, which is a form of gatekeeping (Kraakman 1986; Coffee 2006). We argue that strict liability, if appropriately designed, would incentivize CRAs to supply such certification services as accurately as possible. The idea to introduce strict liability for gatekeepers is not new (Coffee 2004; Partnoy 2006). Importantly, taking into account that the gatekeepers income is very small relative to the investors' losses from underperforming financial assets, these proposals have always capped gatekeepers' liability at a portion of the damages on grounds that a full liability exposure would deter gatekeeping altogether. This problem is germane to that of crushing liability described by Shavell (1980) and Ben-Shahar (2009) among others: if potential injurers are liable for damages in excess to the harm they cause to the society, they may refrain from engaging in valuable activities in the first place. However, despite the possibility to correct for crushing liability, to our knowledge no strict liability proposal has yet been made for rating agencies. This is not without a reason.

The problem of crushing liability is particularly severe for CRAs. The main characteristic of rating agencies is the probabilistic nature of their predictions. To rate a company triple-A is not equal to categorically exclude the possibility of its default; it merely implies a very small probability that default will happen. The logic behind the introduction of a strict liability rule to govern an activity generating losses with a certain probability is that the producer is assumed to be in the best position to insure (or self-insure) against the losses and to raise prices accordingly (Priest 1987). If one tries to apply the same logic to rating agencies, however, three major problems arise.

Firstly, it is possible to insure only against uncorrelated risks. The global financial crisis has shown that, especially in the medium-to-long term, defaults of firms and financial assets can be significantly correlated. For the purpose of this paper, we call these correlations systemic risk. Because strict liability makes the injurer residual risk bearer,

under such regime CRAs would have to pay damages, however limited, stemming from systemic risk, which cannot be insured. Consequently, in order to introduce a workable strict liability rule, the CRAs must be protected against the risk of correlated defaults. We suggest two different ways to deal with this problem, one with respect to the general business risk of corporate bonds and another one, more general, to cope with extraordinary events – like financial crises – which would make strict liability incompatible with the production of ratings, particularly of structured finance products.

Secondly, like other gatekeepers, CRAs cannot face liability for losses significantly larger than the value of their business. Making CRAs pay damages corresponding to the investors' losses from the default of a large rated company would be obviously unreasonable. Because most of these losses would have occurred also in the absence of rating, the expected liability of CRAs could not be possibly compensated by higher fees. Facing such a liability exposure, CRAs would refrain from producing ratings in the first place. Fortunately, it is not necessary to make CRAs liable for the full amount of investors' losses in order for them to have incentives to produce accurate ratings. As suggested by Coffee (2004) for other gatekeepers, it is sufficient to cap the liability at a multiple of their fee income. The problem with this and other similar approaches is the arbitrariness of the multiplier (Haar 2013). We overcome this problem making the multiplier dependent on the probability of default assigned and on the fee received by the CRA. Importantly, under this regime, CRAs cannot make losses unless their predictions are inaccurate.

Thirdly, CRAs cannot be expected to predict default rates without errors. Contrary to a standard assumption in finance, we do not live in a world of perfect foresight. As recently argued by Pacces (2013), it is illusory for the law to police incentives exclusively based on expected values and the underlying probabilities. Our ability to predict the future is limited; so is CRAs' ability to commit to their predictions. Imposing on CRAs a strict liability rigidly dependent on the probabilities they estimate may discourage them from producing ratings in the first place. For this reason, we allow CRAs to reduce their liability exposure by introducing a contractually determined parameter α, which is supposed to account for the uncertainty of the forecasting technology. Through this parameter, the CRAs will be able to prevent crushing liability stemming from the uncertainty of their models, while signaling to the market the degree of confidence in their own estimates.

3. Capped Strict Liability of CRAs: A Numerical Example

As explained in the previous section, imposing unlimited liability on CRAs is not an option. Because the default of any sufficiently large issuer could bankrupt a CRA almost instantly, no ratings would be provided under such regime. However, the characteristics of the market for ratings offer the opportunity to introduce strict liability with a cap on damages based only on objective factors. In the next section, we will show with a formal model that this liability regime is sufficient to deter rating inflation. In this section, we illustrate the intuition of the model with a simple numerical example.

The main task performed by rating agencies is to classify and divide companies in clusters according to their probability of default.[11] To simplify, let us assume that a CRA perfectly knows this probability. If the liability cap is calculated by multiplying the price paid by the issuer times the inverse of the highest probability of default associated with the cluster in which the issuer is included, the liability of the rating agency will depend directly on the extent of rating inflation.

To clarify the idea with a simple example, let us assume there are 100 firms, each one pays $\gamma = 1$ to the CRA for rating, and the cost of rating is zero. Let us also assume that the probability of a default (Pr) is equal to 0.01 for all the firms. If the rating agency correctly estimates the financial stability of the 100 firms, it will include all of them in the same cluster having – we assume – $Pr = 0.01$ as the upper bound. When only one firm effectively goes bankrupt the rating agency will be held liable for $\gamma * 1/Pr = 100$ and will thus make zero profits. It is worth noting that the liability of the CRA is set to 100 independently of the damages stemming from bankruptcy, which could be much higher. However, if the rating agency systematically underestimates the probability of default (that is it inflates the rating), it will bear higher losses. For example, let us assume that all the firms are included in a higher cluster than their creditworthiness would grant, with an assigned probability of $Pr = 0.005$. In this case, if still only one firm goes bankrupt, the liability will be equal to 200, imposing on the CRA a loss of 100.

In this example, we have assumed that CRAs have perfect foresight, that ratings can be produced with zero profits, and that no reputational sanction is attached to rating inflation. In the mathematical model presented in the next section we will relax all these assumptions.

[11] Section 4.3 extends the reasoning to a slightly more detailed discussion of the activity performed by rating agencies.

In concluding this section, it is worth noting that this liability rule compensates investors with a sum of money that is in no way related to the harm they have suffered. However, given that it is nearly impossible to prove CRAs' negligent behavior and the portion of the harm suffered by the investors that is attributable to their conduct, it is hard to determine how much harm rating agencies effectively cause to the market by producing inaccurate ratings. Along with Coffee (2004), we take the approach that in order to improve the incentives of CRAs as gatekeepers, the liability rule should prioritize deterrence over compensation.

4. The Model

Let us define δ_j as a measure of rating inflation (or deflation). With regard to the *jth* cluster of creditworthiness, δ_j is defined as:

$$\delta_j = \frac{m_j - s_j}{m_j} - Pr_j \tag{1}$$

Where the index j varies on the whole set of rating classes, m is the number of firms included by the *nth* CRA in the *jth* class of rating, s_j represents the number of firms included in the *jth* cluster that did not go bankrupt, and Pr_j indicates the default rate for the letter grade associated to the *jth* cluster. In other words, $(m_j - s_j)/m_j$ denotes the ex-post probability of default, whereas Pr_j indicates the ex-ante prediction. Consequently, if CRAs predictions are confirmed ex-post:

$$\frac{m_j - s_j}{m_j} = Pr_j \;;\; \delta_j = 0 \tag{2}$$

Conversely, we formally define rating inflation as:

$$\frac{m_j - s_j}{m_j} > Pr_j \; ; \; \delta_j > 0 \tag{3}$$

The overall level of rating inflation (or deflation) of the *nth* CRA is defined as:

$$\overline{\delta} = \sum_j \delta_j / j \tag{4}$$

We then introduce the parameter β that denotes the difference between the rating assigned to the *ith* firm by the *nth* CRA and the rating potentially assigned to the *ith* by another CRAs. Hence, β measures the level of rating inflation of the *nth* CRA relative to its competitors.

In a perfect market we will assume that the profits of the *nth* CRA are described by the following equation:

$$\Pi_{n,j} = \sum_j \sum_{i=1}^{m_j} \gamma_{i,j} \; + R(\beta, \overline{\delta}) \tag{5}$$

$\gamma_{i,j}$ is the fee collected from each firm net of given rating costs while $R(\beta, \overline{\delta})$ is the reputational effect of CRA's conduct[12]. $R(\beta, \overline{\delta})$ captures the impact of this conduct on future income. For simplicity, we assume a discount factor equal to 1. The reputational effect then depends on the two parameters defined above, namely β and $\overline{\delta}$.

In a perfect market investors will be able to detect any mistake in a CRA's predictions and to punish it with a reputational sanction $R(\beta, \overline{\delta}) < 0$ sufficient to make such mistakes unprofitable. In addition, because there is no market failure, regulation does not need to rely on ratings and there are no regulatory benefits from investing in rated assets. In other words, in a perfect market characterized by perfect foresight, no rating inflation could exist because reputational sanctions are sufficient to prevent opportunistic behavior, regardless of the paying scheme and the liability rule adopted. It follows that in this scenario no liability should be imposed on CRAs.

[12] For the sake of simplicity we assume that CRAs only compete on the number of rated firm, not on the level of the fees. This assumption is without loss of generality, as our results would hold also for variable fees.

We then consider the impact of two market failures, namely the existence of regulatory benefits attached to high ratings and the naivety of some investors. While we are agnostic about the exact impact of each factor, the findings of the literature (Bolton, Freixas and Shapiro 2012; Opp, Opp and Harris 2013) suggest that both m and $R(\beta, \overline{\delta})$ change their shape and their behavior because of them. In this case, the regulator confers an independent value upon high ratings and hence the reputational effect of rating inflation is altered. Under these circumstances it is plausible that conflicting reputational concerns arise. Frenkel (2012) suggests that, especially in concentrated markets, rating agencies facing weak reputational constraints might find it profitable to be lenient and inflate ratings while inducing investors to believe that they are credible. In other words, not only the reputational sanctions might be softened by investors' lack of sophistication, but rating inflation might even be rewarded by institutional investors. As a result, given the existence of regulatory benefits and naïve investors, issuers will be attracted to high ratings regardless of their informative content, and hence m becomes dependent on β and on the size of the regulatory benefits.

Equation (1) thus becomes:

$$\Pi_{n,j} = \sum_{j} \sum_{i=1}^{m(\beta,Rb)_j} \gamma_{i,j} + R\,(\beta, \overline{\delta}, \theta) \qquad (6)$$

Where Rb denotes the regulatory benefits attached to high ratings and θ indicates the share (in value) of naïve investors. Higher values of β and $\overline{\delta}$ result in a higher reputational sanction for the CRA. At the same time, the reputational loss is lower if the value of θ is higher.

Being extremely simple, this description cannot capture the complex nuances that characterize the functioning of CRAs. However, this simple framework is sufficient to include the crucial point made by the literature: given the existing market failures and regulatory distortions, CRAs are able to increase their short-term profits by producing inflated ratings. Under the status quo, CRAs are de facto immune from liability. Therefore, the additional revenues from rating inflation can be larger than the reputational costs to be borne in the future, at least up to a certain level of rating inflation.

Moreover, for individual CRAs, the number of firms to rate depends positively on the level of rating inflation. Because solid firms want to communicate their creditworthiness to the market, some issuers will want to be rated independently of rating inflation. However, another group of issuers will be interested in purchasing a rating *only if* rating inflation is above a certain threshold (for example allowing them to pass the investment grade threshold, which is a condition for investor to enjoy regulatory benefits). Inflating ratings is the only way to attract the issuers of the second group. If this behavior does not sufficiently harm the reputation of the *nth* CRA, rating inflation not only increases short-term profits, but becomes also necessary to survive in the market for ratings. Because the expected liability is nil and the reputational sanctions are not sufficient to support an equilibrium where $\overline{\delta} = 0$, CRAs that do not inflate their ratings will lose customers and short-term profits to their competitors without increasing their future revenues by the same or a higher amount. As a result, all CRAs will inflate ratings to the same extent and the equilibrium will be $\overline{\delta} > 0$ and $\beta = 0$.

Introducing the following liability regime can improve this equilibrium.

4.1. Capped Strict Liability under Simplifying Assumptions

Under the strict liability rule that we propose, the liability of the *nth* CRA for any firm defaulting in the *jth* cluster will be equal to:

$$Ln_j = \sum_{i=1}^{m(\beta,Rb)_j - s_j} \gamma_{i,j} \, / Pr_j \qquad (7)$$

The profits of the *nth* CRA are now equal to:

$$\Pi_{n,j} = \sum_{j} \sum_{i=1}^{m(\beta,Rb)_j} \gamma_{i,j} + R(\beta, \overline{\delta}, \theta) - \sum_{j} \sum_{i=1}^{m(\beta,Rb)_j - s_j} \gamma_{i,j} \, / Pr_j \qquad (8)$$

For the sake of simplicity let us assume that $R(\beta, \overline{\delta}, \theta) = 0$. In other words, for the moment we assume that no reputational sanction is attached to inaccurate ratings.

We obtain for the *jth* cluster:

$$\Pi_{n,j} = \sum_{i=1}^{m(\beta,Rb)_j} \gamma_{i,j} - \sum_{i=1}^{m(\beta,Rb)_j - s_j} \gamma_{i,j} / Pr_j \qquad (9)$$

The ratio $(m(\beta, Rb)_j - s_j)/(m(\beta, Rb)_j$ denotes the share of firms that effectively defaulted. If this ratio is equal to Pr_j then the CRA has correctly estimated the probability of default of the issuer and $\Pi_n = 0$. If the CRA has underestimated the probability of default, which is to say it has inflated the issuer's ratings, then $\Pi_n < 0$. $\Pi_n > 0$ only if $Pr_j > (m(\beta, Rb)_j - s_j)/(m(\beta, Rb)_j$. Hence, facing strict liability according to our model, CRAs will never have any incentive to inflate ratings. To the contrary, the optimal strategy for them would be to award always a probability of default equal to 1. This extreme case of rating deflation is purely theoretical, because obviously no issuer will ever be interested in purchasing such a rating. Actually, also because highly rated assets bring about regulatory benefits to regulated investors, issuers will have an interest to receive a rating that is as high as possible.

Issuers, CRAs and regulated investors have normally an information advantage compared to regulatory authorities and courts. The question is how to induce the market for ratings to reveal information efficiently. Our strategy is to create, by imposing an appropriate strict liability on CRAs, opposing interests for CRAs, issuers and investors. More specifically, the CRAs will prefer to supply lower ratings in order to reduce their expected liability, whereas issuers and regulated investors will prefer higher ratings. The ratings produced in such a market are going to reflect valuable information about the creditworthiness of issuers and their bonds. In fact, this is the only way in which gains from trade can be generated after the profits from misrating are disallowed by a capped strict liability rule. This outcome will ultimately benefit financial regulators and the society at large.

In every market the opposing interests of sellers and buyers lead to an equilibrium price that, absent market failures, is considered optimal. To re-create such equilibrium in the market for ratings we make sure that issuers and regulated investors, on one side, and

CRAs, on the other side, have opposite interests. This has also important dynamic implications. Under the status quo, increasing competition between CRAs would only worsen the problem of rating shopping (Becker and Milbourn 2011; Bolton, Freixas and Shapiro 2012). This circumstance rules out the most straightforward strategy to improve the efficiency of ratings, namely increasing competition. Competition could again be valuable in the market for ratings after imposing strict liability on CRAs. In the presence of a capped strict liability regime like the one that we have designed, more actual and potential competition between CRAs can be expected to lead to more innovation in forecasting techniques rather than to more rating inflation.[13]

4.2. Capped Strict Liability with Imperfect Foresight and Reputational Sanctions

Under our liability rule, four different conditions have to be fulfilled for an efficient market for ratings to emerge: *(i)* $\Pi = 0$ is considered a satisfying equilibrium; (*ii*) rating agencies know the true probability of default; (*iii*) $R(\beta, \overline{\delta}, \theta) = 0$ and (*iv*) firms defaults are uncorrelated.

With respect to (*i*), to use γ as the relevant base for our liability rule implies that the profits of CRAs, given accurate ratings, are set to zero. They become negative only in the presence of rating inflation, which under the assumption of perfect foresight is sufficient to guarantee rating accuracy. The condition $\Pi = 0$ is reminiscent of the absence of economic profits under perfect competition and is not particularly restrictive. As mentioned in the previous section, this equilibrium cannot be improved by exaggerating the probability of default (rating deflation) because at some point this will drive the number of rated firms to zero. This scenario is not particularly interesting for policymaking; therefore, we are not considering it in this paper.

More importantly, even under ideal incentives, the CRAs will be prone to make mistakes, violating condition (*ii*). In fact, condition (*ii*) is never true – we do not live in a world of perfect foresight. In addition, the assumption *(iii)* – namely that R = 0 – should be relaxed too in order to take into account the effects of reputation and, more in general, all the factors affecting the future income of CRAs. Finally, condition (*iv)* concerns systemic risk as a source of crushing liability. We will deal with this problem in section 5.

[13] We expand on this point below, in section 6.

To address (*ii*) and (*iii*) we introduce the parameter $0 < \alpha < 1$. α limits the expected liability of CRAs. After the introduction of α, the profit of the *nth* CRA will be equal to:

$$\Pi_{n,j} = \sum_{j} \sum_{i=1}^{m(\beta,Rb)_j} \gamma_{i,j} + R(\beta, \overline{\delta}, \theta) - \sum_{j=D} \sum_{i=1}^{m(\beta,Rb)_j - s_j} \gamma_{i,j} * \alpha / Pr_j \qquad (10)$$

Where α denotes the fraction of $\gamma_{i,j}$ that is considered to calculate the expected liability. The smaller α, the more mistakes CRAs are allowed to make without suffering losses (and the more economic profits they can make if their ratings are correct). In other words, this scenario lies between two extremes: a perfect market where ratings are efficiently policed by reputational concerns; and the stylized market described in section 4.1 in which $R(\beta, \overline{\delta}, \theta)$ = 0 and rating agencies face liability whenever a rated issuer defaults. In the former case the optimal α (let that be α^*) is equal to 0, in the latter it is equal to 1. Because, as shown by the literature on CRAs, the reputational sanction is neither optimal nor is it totally absent, α^* will lie between the two extremes.

Identifying such an optimal value might seem attractive, but we believe that this would be an almost impossible task. A benevolent and omniscient regulator could identify the optimal value of α for any transaction and at any moment in time. However, an omniscient regulator would also know the correct rating for any issuer and financial asset and thus the whole problem of accurate ratings would simply not arise. On the contrary, regulators neither possess unlimited information nor can they be expected to be always benevolent. It seems extremely difficult that a public authority can adequately manipulate α in order to guarantee that CRAs earn enough to stay in business without being tempted to inflate their ratings. In order to determine α^*, it would be necessary to know the value of the parameters $\beta, \overline{\delta}, \theta$, the shape of the functions $R(\beta, \overline{\delta}, \theta)$ and m, and the level of accuracy of the available forecasting technology.

In more qualitative terms, we argue that the simultaneous presence of regulatory benefits, naïve investors, and imperfect forecasting techniques has affected the market for ratings in a very complex way. In our view, re-creating opposing interests between supply and demand for ratings is a better strategy than attempting to correct the above reasons for market failure via detailed regulations. Given the existing market failures, we expressly refuse to make further assumptions about the shape of $R(\beta, \overline{\delta}, \theta)$. We prefer a solution that relies on market mechanisms to determine α, based on the market players' knowledge of

the parameters determining the size of the reputational sanction $R(\beta, \overline{\delta}, \theta)$. Obviously, the higher is α, the more CRAs will be credible because they are punished if they inflate their ratings. However, the expected liability may be too high to sustain a market for ratings given the existing forecasting technology. A lower α, on the other hand, is good to keep CRAs in business, but might be insufficient to cope with the problem of rating inflation given the shape of $R(\beta, \overline{\delta}, \theta)$.

The alternative to choosing α by regulation is to let α be determined contractually. We take this approach by allowing CRAs to announce to the market (that is, to the investors) how much they are committing to a certain rating with their choice of α. This approach copes with an important shortcoming of imposing strict liability on CRAs. CRAs have often stated that their predictions are ordinal in nature, not cardinal. The liability rule that we propose requires that all CRAs be compelled to publish the specific range of probability of default associated to a certain rating, and particularly to connect the upper bound of this range to their expected liability. In a sense, this implies forcing CRAs to produce ratings as a cardinal measure. Although this increases transparency, it would also place on CRAs a burden that they might be unwilling to bear. If the value of α is determined by a regulatory authority, there is the concrete risk that this burden becomes excessive. As we mentioned, regulators are not omniscient. Neither are CRAs. Imposing on CRAs a given α means committing them to a given level of confidence in their own probability estimates. CRAs that find such a level of confidence excessive may simply decide to exit the market.

Conversely, if the rating agencies are allowed to decide how much to "bet" on a certain rating, they will be able to take into account the unavoidable uncertainty surrounding predictions of the future and the possibility of mistakes or imperfection in their models. This solution has a number of advantages. Firstly, it introduces a commitment device to improve the functioning of the market for ratings. This device is a varying degree of liability exposure, which CRAs can choose freely so long as this choice allows them to produce ratings valued by investors. Secondly, because the CRAs know better than anybody else how accurate their forecasting models are in predicting future defaults, they can choose the level of commitment that is sufficient to keep them in the business thus preventing strict liability from becoming crushing (Shavell 1980).

The key feature of α is its contractibility. Being a commitment device supported by an enforceable strict liability rule, α can be as low as to keep CRAs in business and as high as

to make ratings informative for investors including the naïve ones.[14] In other words, α allows contracting on unobservable parameters like the determinants of $R(\beta, \overline{\delta}, \theta)$ and the uncertainty of forecasting models. In the absence of regulatory distortions, competition in the provision of certification services to issuers will always make sure that α is the efficient outcome of the opposing interests of CRAs and investors. Moreover, because CRAs will compete on α, this mechanism also provides incentives to improve the forecasting technology over time. Only the presence of regulatory benefits from high ratings makes this market approach unviable, because such benefits could be so high as to offset all the negative determinants of $R(\beta, \overline{\delta}, \theta)$. When this is the case, the regulatory benefits can sustain a market for ratings also with α artificially low (or even zero).

If α is contractually determined, financial regulation cannot allow whatever rating to have regulatory relevance. More precisely, besides requiring a high rating for investors to enjoy regulatory benefits, regulation should also impose that α chosen by the CRA producing the rating is above a specific threshold. Under such arrangement, rating agencies would not merely claim that a firm deserves a high rating, but they would have to put their money where their mouth is in order to be credible. At the same time, by deciding exactly how much to expose themselves to liability, CRAs can prevent the risk that an excessively zealous regulator forces them to carry an excessive burden − at the end CRAs are not obliged to produce rating relevant to regulation. It is important to note that CRAs are not forced to adopt any particular value of α. In theory, they could simply decide to shield themselves from any liability claim if that was acceptable for issuers and investors. However, if CRAs want their ratings to have a regulatory value, they should be the first to show reliance in their own predictions by complying with a minimum value of α established by regulation.

4.3. Extending the Model: Loss Given Default

In certain cases, especially for corporate bonds, ratings are not only an indicator of the probability of default, but also include an estimate of the loss given default (LGD).

In this section our model is adapted to take into account the LGD as well as any other quantitative aspect that CRAs might consider to produce a rating. Once again, for the sake

[14] We assume that no investor is so naïve to be unable to rank commitments to liability exposure based on $0 < \alpha < 1$.

of simplicity, we will refer to equation (5) under the assumption of perfect foresight. To take into account the LGD, equation (5) should be modified in the following way:

$$\Pi_{n,j} = \sum_{i=1}^{m(\beta,Rb)_j} \gamma_{i,j} - \sum_{i=1}^{m(\beta,Rb)_j - s_j} \gamma_{i,j} / P_{r_j} * LGD_{r_i} / LGD_{p_i} \qquad (11)$$

LGD_r represents the LGD effectively observed whereas LGD_p represents the predicted LGD. Similarly to our previous discussion on the probability of default, if $LGD_r > LGD_p$ then the expected profits will decrease. If $LGD_r = LGD_p$ the expected profits will not be altered by liability. Lastly, for $LGD_r < LGD_p$, Π_n would theoretically increase, but as we explained for the probability of default, a scenario in which CRAs systematically underestimate creditworthiness is not very realistic because, at some point, issuers will simply stop buying its ratings.

This simple extension shows that our liability rule could be applied, with an identical logic, to any quantitative factor employed by rating agencies for the production of their assessment.

5. Systemic Risk

To avoid that strict liability becomes crushing, it is necessary to protect CRAs from systemic risk, which may result in correlated defaults. Correlated defaults are problematic both because they undermine the ex-post accuracy of CRAs' estimates and because they are a risk that cannot be insured (or self-insured) by definition. While we account for the fallacies of forecasting models through the choice of α, our strict liability rule still makes CRAs residual risk bearer for the portion of damage compensation triggered by the default of a rated issuer or bond. Therefore, apart from the uninteresting case in which α is set to 0 (or precisely because we want to avoid this outcome), it is important to make sure that CRAs do not face liability when defaults depend on systemic risk rather than on the individual circumstances of the issuer or of the bond that ratings are supposed to assess with a varying degree of precision (α).

Unfortunately, there is no unique way to cope with this problem. As shown by the recent literature (Coval, Jurek and Stafford 2009 a,b: Rablen 2013) rating structured finance products differs from rating traditional corporate bonds precisely because of their different exposure to systemic risk. As we are going to show, corporate bonds are rather insensitive to fluctuations of economic output in the short term. This offers a straightforward way to deal with systemic risk: the strict liability of CRAs should be limited to the short term. However, structured finance products are very different from corporate bonds in this respect because their defaults can be highly correlated also in the short term. To be sure, contrary to traditional corporate bonds whose credit risk mainly depends on firm-specific characteristics, structured finance products behave like economic catastrophe bonds (Coval, Jurek and Stafford 2009a) concentrating defaults in the worst states of the economy as a whole. This extreme sensitivity of structured finance to systemic risk is a problem that cannot be ameliorated limiting the CRAs' liability to the short term. Therefore, we reserve this approach exclusively to corporate bonds. As far as structured finance products are concerned, addressing systemic risk requires a modification of our strict liability regime. We present these two approaches in turn.

5.1. Short-Term Liability for Rating Corporate Bonds

Predictions can be medium-to-long term or short term. We consider three months a typical short-term horizon because this is usually the timeframe (the so-called "watchlist") in which CRAs review their assessment and decide whether to maintain or downgrade a certain rating (Bannier and Hirsh 2010). The rating of corporate bonds mainly depends on the probability that their issuers – typically business enterprises – go bankrupt.[15] While medium-to-long term predictions in this respect seem to be greatly affected by systemic risk, short-term predictions present this problem in an attenuated form. If the focus is a sufficiently short time horizon, there is no reason to expect that the correlation between business issuers going bankrupt will be significantly positive. This seems to hold true even in times of aggregate economic distress. For instance, the data from the Quarterly U.S. Business Bankruptcies show that even during a crisis as violent as the global financial crisis, bankruptcies have taken a certain time to propagate.

[15] For simplicity we do not include another important determinant, namely the Loss Given Default. See section 4.3.

Figure 1 illustrates this well. It can be noticed that, although the increase in the frequency of bankruptcies between 2006 and 2009 was significant, the short-term fluctuations were not particularly violent.

[Figure 1 about Here]

The point is illustrated even more clearly by the contrast between Figure 2 and Figure 3. If we look at a period of one calendar year, the percent change in the number of bankruptcies is dramatic, reaching peaks of 44% and 54% respectively in 2007 and 2008. On the contrary, if we consider a shorter horizon, for instance a quarter, the percent changes are much smaller. These changes are often below the ten percent threshold, and are never above 19%. Without pretense to discuss thoroughly the impact of systemic crises on bankruptcy rates, we want to emphasize that these data suggest that firm defaults can be indeed correlated; but economic crises, however severe, do not spread instantly across issuers.

[Figure 2 about Here]

[Figure 3 about Here]

Based on this observation, the strict liability faced by CRAs rating corporate bonds should have an expiration date. Rating agencies would be strictly liable only if the issuer goes bankrupt shortly after the rating has been issued or confirmed. If the definition of short term coincides, as we suggest, with the typical interval in which CRAs review their ratings, CRAs will have the opportunity to revise their ratings when changed circumstances call for a different assessment. If an aggregate shock takes longer than three months to alter the frequency of defaults, CRAs will avoid liability just by adjusting their ratings to the new environment when the revisions come due. At the same time, liability cannot be avoided simply by downgrading firms that suddenly turn out to be riskier than originally foreseen. Once a rating is given or is confirmed, it will commit the CRA for three months in a proportion corresponding to the choice of α as defined in section 4.2. After the expiration date, the standard negligence rule could be put back in place, which is another way to say that CRAs would face no liability, as is currently the case.

5.2. Postponed Liability for Rating Structured Finance Products

Although limiting liability to the short term offers CRAs an effective protection against systemic risk in the case of corporate bonds, this solution may not be sufficient for structured finance and, more in general, whenever defaults can be positively correlated also in the short term. Under these circumstances, the liability of CRAs simply needs to be excluded if defaults depend on systemic risk. In order to achieve this result, we must depart from the traditional tort law approach. In this perspective, however, it is still possible to imagine an incentive scheme grounded on the same model presented in the previous section, with the modifications below.

This system would work as follows. A public authority records the rating issued by the CRAs, the fees they receive, and the actual frequency of defaults of each structured finance product. Using α/Pr as a multiplier, the regulator calculates the potential liability that each CRA has to face for each default. CRAs should still be allowed to choose α as in the strict liability regime designed before. However, CRAs will not be asked to pay damage compensation whenever a structured finance asset defaults. Only after a certain time interval, say one year, the public authority will verify the overall accuracy of a CRA's predictions, which, in turn, will determine whether the CRA in question is to face liability for the assets that defaulted in the previous year. For example, let us consider the cluster BBB- (Baa3 using Fitch scale). The historical, annualized range of probability of default associated with this cluster is 0.025 - 0.032.[16] If, during the time interval considered, less than 0.032 of the assets included in the cluster have defaulted, then no compensation will be due. Conversely, liability will be triggered if the quality of ratings has been below the relevant threshold. In other words, the payment will be due only if more than 0.032 of the assets included in the cluster BBB- has defaulted.

Postponing the imposition of the monetary sanction allows making liability conditional on the failure of CRAs to predict default over a sufficiently large number of observations. This approach has two advantages in coping with systemic risk. Firstly, if the predictions of rating agencies turn out not to be inflated over the relevant timeframe, their profits will not be affected by the defaults occurring within their range of predictions because they will simply face no liability for those defaults. Compared to the strict liability solution, this mechanism tempers the over-deterrence stemming from the uninsurability of systemic risk. However, CRAs would still be liable to pay damages when the frequency of default in a given time interval exceeds the highest probability of default in the relevant class of rating.

[16] This example is based on Fitch's historical (annualized) default experience. See Coval, Jurek and Stafford (2009b: 9)

This effect is desirable to police rating inflation; but it also leaves CRAs exposed to systemic risk, particularly in those scenarios of 'economic catastrophe' where structured finance assets tend to experience extraordinary rates of defaults. Financial crises are a case in point.

Fortunately, postponing the imposition of liability has a second advantage in coping with systemic risk. The timeframe for assessing the accuracy of CRAs' prediction could be made long enough to absorb the violent fluctuations in the default of structured finance products depending on a financial crisis. Obviously, for this purpose, the length of the interval is crucial. Whereas a one-year period could be sufficient to assess the accuracy of CRAs' ratings of structured finance in normal times, this might be just too short a time to compensate the sudden spikes in defaults coming along with a financial crisis. For this reason, we advocate the introduction of a double layer of protection for the rating of structured finance products. At a first stage the ex-ante predictions of rating agency are compared with the ex-post default rates during the year in question. As stated above, if the predictions are accurate over one year, no liability will be imposed on rating agencies. Conversely, if the CRA has underestimated the number of defaults over one year, the public authority could decide *on an exceptional basis* to impose liability on the additional condition that ratings were inflated also over a longer time horizon. Importantly, in order to protect CRAs from systemic risk, the relevant timeframe can be extended backward, not forward. If, because of a financial crisis, structured finance products have experienced extraordinary rates of defaults in a year, it will take many years before the situation returns to normality and even longer before the shock can be absorbed by the data.

Let us illustrate our solution with a simple numerical example. Assume that, for instance over the past five years, a rating agency has predicted for a given class of structured finance products the expected number of defaults (ED) indicated in the table below. Let also the actual number of defaults (ND) be as reported in the following table.

Year	***1***	***2***	***3***	***4***	***5***	∑
ED	10	9	8	7	6	40
ND	9	8	7	6	10	40

Only in the fifth year the rating agency has underestimated the number of defaults. Therefore, under the normal rule, the CRA should be liable to pay a compensation based on the α/Pr multiplier. However, the public authority might exceptionally determine that a spike from 6 to 10 defaults from one year to another is a consequence of systemic risk and hence it might extend the assessment interval. For example, regulation may provide that in such situations the assessment interval could be extended up to the average maturity of the structured finance products in question. If we posit, as in the numerical example, that the average maturity is five years, the latter becomes the relevant timeframe to determine whether the CRA is liable. As the example shows, over a five-year period ED is equal to ND (40 defaults) and therefore, despite the spike in the number of defaults in the fifth year, the CRA will face no liability.

This solution would reward the CRAs who were more conservative in their ratings the years preceding a financial crisis, as those CRAs could count on historical frequencies of default below the maximum PD associated with the relevant letter grade. This effect is countercyclical, namely it counters, however little, the formation of asset bubbles without standing in the way of a recovery of credit.[17] Although this solution ultimately relies on the ex-post discretion of public authorities to cope with systemic risk – which we assume to be unpredictable – it is worth noting that this discretion is essentially limited to the declaration of a status of financial crisis; all of the other consequences on CRAs' liability should be predetermined by regulation.

6. The Virtues of Capped Strict Liability

It is worthwhile to briefly highlight the benefits of the approach presented in this paper. In the first place, the liability rule that we propose connects CRAs' expected profits to the quality of their ratings, thus inducing them to put their money where their mouth is. In fact, by tying the expected liability to the rating assigned (and to the level of commitment accepted), the CRAs' profits will depend on the quality of their predictions. It follows that the problem of rating shopping is addressed implicitly, as any CRA that produces overoptimistic ratings to attract more issuers will be forced to face higher liability.

[17] The countercyclical property of our solution is especially relevant given that ratings quality has been shown to be lower during booms (Bar-Isaac and Shapiro 2013).

Secondly, our proposal introduces a damage cap based on objective factors. The cap has the important virtue to prevent over-deterrence of rating activity. At the same time, this approach eliminates almost any discretion on the side of regulators and courts. The only exception is the declaration of financial crisis status necessary to offer CRAs rating structured finance products a second layer of protection against systemic risk. Otherwise, the rule we advocate carries sizeable savings in terms of administrative costs. There will be no need to scrutinize the behavior of CRAs or to establish complex standards of care in order to prevent them from producing inflated ratings. Moreover, courts will not have to quantify the portion of damages attributable to the conduct of CRAs. Determining that an issuer or a bond have defaulted and multiplying the price by the probability of default associated with a given rating are (quasi-)automatic and (quasi-)costless tasks. The risk of litigation errors, frivolous litigation, and opportunistic settlements usually created by strict liability (Coffee 2004) will all be ruled out.

Finally, the incentive scheme described above ties the income earned by CRAs to the quality of their forecasting techniques, thus creating the right incentives also from a dynamic perspective. To understand this point, let us assume that three rating agencies exist: A, B, and C. A and B have developed state-of-the-art forecasting models; thus they are able to assess with greater accuracy than C the issuer's probability of default. Let us also assume that firms are divided in two groups, X and Y, which respectively have a low and a high probability of default. Under these assumptions C will not be able to distinguish between X and Y and will therefore be forced to assign an average probability of default. Good issuers, however, could obtain better ratings from A and B because these rating agencies are able to better assess their creditworthiness. As a consequence, issuers belonging to the cluster X will switch to the two CRAs that are able to assign them the rating they deserve. The more good firms switch to A and B, the higher will be the average level of risk of the pool of firms rated by C. In the end, all the good firms that have a low probability of default will be rated by A and B, and the firms with a high probability of default will be indifferent between A, B and C. As, in the real world, the probability of default of rated firms approaches a continuous function, the only competitive equilibrium is one where every firm opt for A or B, unless they are so risky to be indifferent between A, B, and C. In this case, however, the rating would have hardly any added value for the issuer and C would have to exit the market for ratings.

An identical reasoning applies to the parameter α when CRAs choose freely how much to expose themselves to liability. In fact, CRAs that can offer predictions which are more accurate will be able to determine with higher precision when they can expose themselves to a higher liability. It is obvious that good firms will have every incentive to hire the CRA that can adopt a higher value of α, both because this implies a higher commitment to rating accuracy and because – as we argued – a relatively high α should be a precondition for ratings to deliver regulatory benefits. For analogy with the mechanism described above, a high α will emerge as a result of CRAs' competition on the quality of forecasting techniques.

7. Conclusion

There has been an enormous debate both at the political and at the academic level on how to improve CRAs' incentives to produce accurate ratings. In this paper, we propose the introduction of an expiring, capped strict liability rule with a contractual component. A damage cap based on objective factors is introduced in order to avoid crushing liability, whereas the expiration date is needed to shield CRAs from systemic risk whenever – as in the case of corporate bonds – defaults are largely uncorrelated in the short term. Furthermore, CRAs are allowed to determine contractually at what level they want to commit to their predictions. Importantly, no liability is imposed on them, unless they want their ratings to have regulatory relevance. Finally, in order to protect CRAs from systemic risk also when defaults can be correlated in the short term (as in the case of structured finance), we propose a departure from the traditional tort law approach. By delaying the compensation until after few defaults have occurred, CRAs may be punished only when their predictions are proven to be inaccurate over a sufficiently large time interval.

References

[1] Baklayar, Inna, and Koresh Galil. 2011. "Rating Shopping and Rating Inflation: Empirical Evidence from Israel." *Available at http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1903827*

[2] Bar-Isaac, Heski, and Joel Shapiro. 2013. "Ratings Quality over the Business Cycle." *Journal of Financial Economics* 108: 62-78.

[3] Becker, Bo, and Todd Milbourn. 2011. "How Did Increased Competition Affect Credit Ratings?" *Journal of Financial Economics* 101.3: 493-514.

[4] Ben-Shahar, Omri. 2009. "Causation and Foreseeability." 83-108 in *Tort Law and Economics*, edited by Michael Faure. Cheltenham: Edward Elgar.

[5] Bolton, Patrick, Freixas Xavier and Joel Shapiro. 2012. "The Credit Ratings Game." *Journal of Finance* 67.1: 85-112.

[6] Calabresi, Guido. 1970. *The Cost of Accidents: A Legal and Economic Analysis.* New Haven: Yale University Press.

[7] Calomiris, Charles W. 2009a. "A Recipe for Ratings Reform." *The Economists' Voice* 6.11: 1-5.

[8] Calomiris, Charles W. 2009b. "The Subprime Turmoil: What's Old, What's New, and What's Next." *The Journal of Structured Finance* 15.1: 6-52.

[9] Coase, Ronald H. 1960. "Problem of Social Cost." *The Journal of Law and Economics* 3: 1-44.

[10] Coffee Jr, John C. 2004. "Gatekeeper Failure and Reform: The Challenge of Fashioning Relevant Reforms." *Boston University Law Review* 84: 301-364.

[11] Coffee Jr, John C. 2006. *Gatekeepers: The Professions and Corporate Governance.* Oxford and New York: Oxford University Press.

[12] Coffee Jr, John C. 2010. "Ratings Reform: The Good, the Bad and the Ugly." *Harvard Business Law Review* 1: 232-276.

[13] Cooter, Robert D., and Thomas S. Ulen. 2011. *Law and Economics*. 6th Edition. Boston, MA: Addison-Wesley.

[14] Coval, Joshua D., Jakub W. Jurek, and Erik Stafford. 2009a. "Economic Catastrophe Bonds." *The American Economic Review* 99.3: 628-666.

[15] Coval, Joshua, Jakub Jurek, and Erik Stafford. 2009b. "The Economics of Structured Finance." *The Journal of Economic Perspectives* 23.1: 3-26.

[16] Deats, Caleb. 2010. "Talk that isn't cheap: Does the First Amendment Protect Credit Rating Agencies' Faulty Methodologies from Regulation." *Columbia Law Review* 110: 1818-1864.

[17] Flannery, Mark J., Joel F. Houston, and Frank Partnoy. 2010. "Credit Default Swap Spreads As Viable Substitutes for Credit Ratings." *University of Pennsylvania Law Review* 158.7: 2085-2123.

[18] Frenkel, Sivan. 2012. "Repeated Interaction and Rating Inflation: A Model of Double Reputation." Unpublished working paper. Hebrew University.

[19] Gorton, Gary and Guillermo Ordoñez. 2014. “Collateral Crises.” *American Economic Review* 104.2: 343-378.

[20] Haar, Brigitte. 2013. “Civil Liability of Credit Rating Agencies after CRA 3-Regulatory All-or-Nothing Approaches between Immunity and Over-Deterrence.” *University of Oslo Faculty of Law Research Paper,* February.

[21] Heremans, Dirk, and Alessio M. Pacces. 2012. “Regulation of Banking and Financial Markets.” 558-606 in *Regulation and Economics*, edited by Alessio M. Pacces and Roger J. Van den Bergh. Cheltenham: Edward Elgar.

[22] Kraakman, Reinier H. 1986. “Gatekeepers: the anatomy of a third-party enforcement strategy.” *Journal of Law Economics & Organizations* 2: 53-104.

[23] Krugman, Paul. 2010. “Berating the raters.” *New York Times A 23:* 144-149.

[24] Listokin, Yair. 2010. “If You Misrate, Then You Lose: Improving Credit Rating Accuracy Through Incentive Compensation.” *Yale Journal on Regulation* 27: 91-113.

[25] Mathis, Jerome, James McAndrews, and Jean-Charles Rochet. 2009. “Rating the Raters: Are Reputation Concerns Powerful Enough to Discipline Rating Agencies?” *Journal of Monetary Economics* 56.5: 657-674.

[26] Opp, Christian, Opp Marcus, and Milton Harris. 2013. “Rating Agencies in the Face of Regulation” *Journal of Financial Economics* 108: 46-61.

[27] Pacces, Alessio M. 2013. *The Future in Law and Finance*. Den Haag: Eleven International Publishing.

[28] Pagano, Marco, and Paolo Volpin. 2010. “Credit Ratings Failures and Policy Options.” *Economic Policy 25*.62: 401-431.

[29] Park, Sun Young. 2011. “The Size of the Subprime Shock.” Unpublished manuscript, Korea Advanced Institute of Science and Technology.

[30] Partnoy, Frank. 1999. “The Siskel and Ebert of Financial markets: Two Thumbs Down for the Credit Rating Agencies.” *Washington University Law Quarterly* 77: 619-712.

[31] Partnoy, Frank. 2006. “How and why credit rating agencies are not like other gatekeepers.” 59-102 in *Financial Gatekeepers: Can They Protect Investors?*, edited by Yasuyuki Fuchita and Robert E. Litan. Baltimore: Brookings Institution Press and Nomura Institute of Capital Markets Research.

[32] Partnoy, Frank. 2010. “Rethinking Regulation of Credit-Rating Agencies: An Institutional Investor Perspective.” *Journal of International Banking Law* 25.4: 188-199.

[33] Priest, George L. 1987. “The Current Insurance Crisis And Modern Tort Law.” *The Yale Law Journal* 96.7: 1521-1590.

[34] Rablen, Matthew D. 2013. “Divergence in Credit Ratings.” *Finance Research Letters* 10.1: 12-16.

[35] Sangiorgi, Francesco, Jonathan Sokobin, and Chester Spatt. 2009. “Credit-Rating Shopping, Selection, and the Equilibrium Structure of Ratings.” Manuscript, Carnegie-Mellon University.

[36] Shavell, Steven. 1980. “Strict Liability Versus Negligence.” *The Journal of Legal Studies* 9.1: 1-25.

[37] Shavell, Steven. 1987. *Economic Analysis of Accident Law*. Cambridge, MA: Harvard University Press.

[38] Skreta, Vasiliki, and Laura Veldkamp. 2009. “Ratings Shopping and Asset Complexity: A Theory of Ratings Inflation.” *Journal of Monetary Economics* 56.5: 678-695.

[39] Strobl, Gunter, and Han Xia. 2012. “The Issuer-Pays Rating Model and Ratings Inflation: Evidence from Corporate Credit Ratings.” Unpublished working paper.

[40] White, Lawrence J. 2010. “Markets: The Credit Rating Agencies.” *The Journal of Economic Perspectives* 24.2: 211-226.

Figures

FIGURE 1: BANKRUPTCIES OF BUSINESS FIRMS IN THE U.S. (IN THOUSANDS)

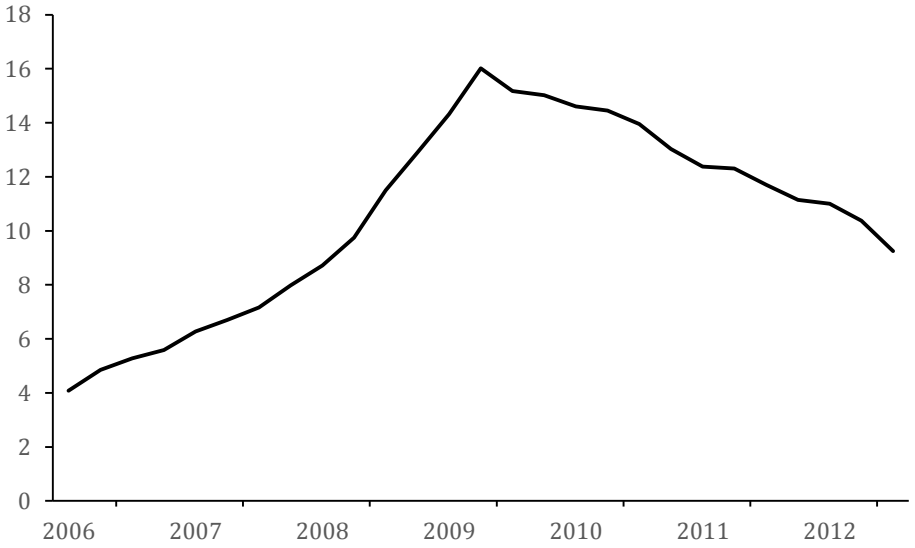

SOURCE: AMERICAN BANKRUPTCY INSTITUTE (www.abiworld.org)

FIGURE 2: YEARLY PERCENT CHANGE IN BANKRUPTCIES OF BUSINESS FIRMS IN THE U.S.

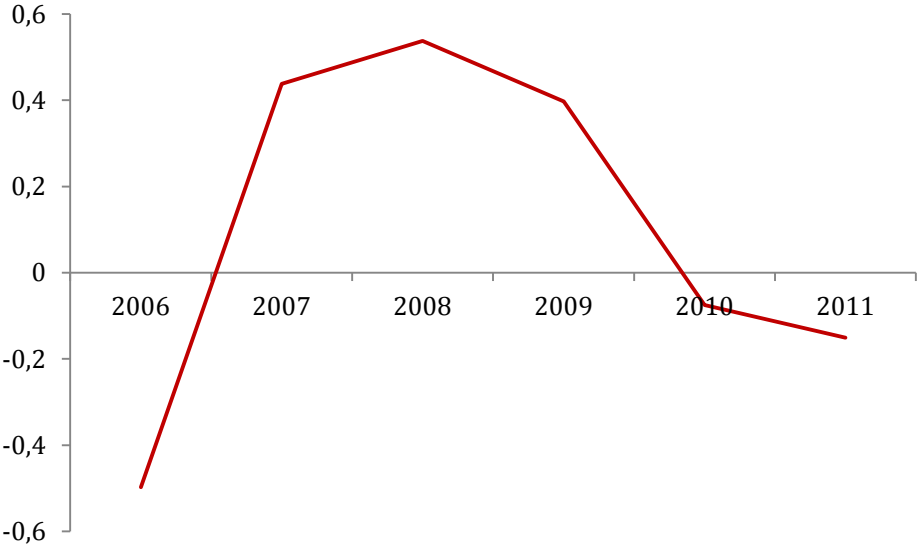

SOURCE: AMERICAN BANKRUPTCY INSTITUTE (www.abiworld.org)

FIGURE 3: QUARTERLY PERCENT CHANGE IN BANKRUPTCIES OF BUSINESS FIRMS IN THE U.S.

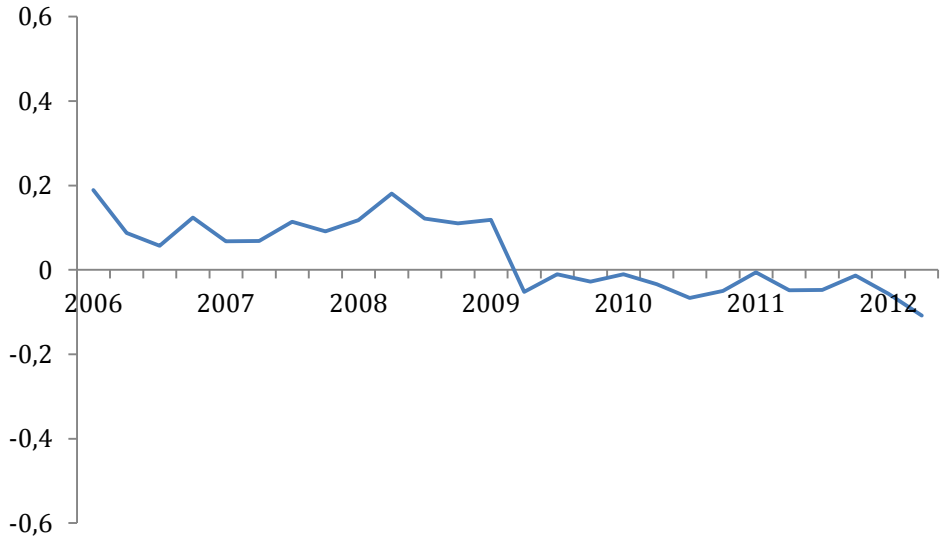

SOURCE: AMERICAN BANKRUPTCY INSTITUTE (www.abiworld.org)

about ECGI

The European Corporate Governance Institute has been established to improve *corporate governance through fostering independent scientific research and related activities.*

The ECGI produces and disseminates high quality research while remaining close to the concerns and interests of corporate, financial and public policy makers. It draws on the expertise of scholars from numerous countries and bring together a critical mass of expertise and interest to bear on this important subject.

The views expressed in this working paper are those of the authors, not those of the ECGI or its members.

Electronic Access to the Working Paper Series

The full set of ECGI working papers can be accessed through the Institute's Web-site (www.ecgi.org/wp) or SSRN:

Finance Paper Series	http://www.ssrn.com/link/ECGI-Finance.html
Law Paper Series	http://www.ssrn.com/link/ECGI-Law.html