Re-Examining the Role of Rating Agencies: Lessons from Structured Finance

Joshua Coval, Jakub Jurek, & Erik Stafford

January 2008

Credit markets have traditionally played the dual-role of serving as a source of funding for governments, municipalities, households, and corporations alike, and as a mechanism for redistributing risks across market participants who are best suited to bear them. Although these roles have mainly been immutable across time, the variety of credit instruments available to perform these roles has proliferated and includes everything from simple coupon bonds to bonds backed by the cash flows from large pools of related assets (e.g. asset-backed securities). More recently, this array of securities was expanded even further by the advent of structured finance, which has enabled originators to manufacture highly-rated securities out of other lower-rated financial products. The prototypical structured finance product is the collateralized debt obligation (CDO), whereby a large portfolio of credit sensitive securities is formed and then a new capital structure is issued against the underlying portfolio. The most senior claims in this new capital structure are able to obtain ratings that are higher than the average rating of the underlying portfolio, as they are protected from initial losses by the junior claims. While many investors have come to view the resulting CDO claims as comparable to traditional single-name corporate bonds – they have finite maturities, typically bear coupons and are rated by credit rating agencies, such as Moody's, S&P and Fitch – this paper argues that they materially differ from the perspective of both their cash flows and systematic risk exposures, and thus are fundamentally distinct financial products.

A tremendous amount of structured finance securities have been issued over the course of the last few years. It is estimated that the outstanding notional value of US structured products had reached $9 trillion in 2007,[1] equaling that of the US Treasury market. The recent explosion

[1] Financial Times, "Failing grades? Why regulators fear credit ratings may be out of their depth" May 17, 2007.

of structured finance securities has tested the institutional forms that have evolved in credit markets over the course of the past century, in particular, the central role of credit rating agencies. Although historical evidence is mixed regarding the actual value of ratings, many market participants have come to rely on them in numerous ways. Issuers use ratings to alleviate information asymmetries and thereby expand the pool of potential buyers of their securities. Investors use ratings as a basis for pricing and risk assessment. Regulators rely on ratings in setting investment guidelines and determining capital requirements. This reliance on ratings has continued in structured finance markets, where the added complexity of the securities being manufactured by Wall Street banks created a natural role for the services of an independent third-party monitor, such as a credit rating agency. In effect, by evolving as a rated market, the investor base remained largely the same as that for traditional credit securities, enabling the market to operate within existing regulations.

A defining feature of recent structured finance activities has been the large quantities of triple-A rated securities manufactured. As of mid-2007 there were 37,000 structured finance issues in the U.S. alone with the top rating, accounting for nearly $5 trillion in notional value.[2] According to Fitch Ratings (2007), roughly 60% of all global structured products were AAA-rated, in contrast to less that 1% of the corporate and financial issues. Since many large institutional investors, such as pension funds and insurance companies, are restricted to owning highly-rated securities, many have argued that the growth of structured finance reflects its ability to "complete the market" for these assets. In principle, these products can promote better risk sharing if all parties fully understand how structured securities work (DeMarzo (2005)), by allowing cautious investors to obtain low risk exposures by investing in senior claims.

There may be a darker side to these products as well. The complex nature of structured finance products allowed credit rating agencies to transport their reputation from the single-name market to the structured finance market. It remains an open question as to whether expertise in

[2] Financial Times, "Top rating proving crucial to structured finance sector" May 17, 2007.

rating corporate bonds, and even ratings themselves, are relevant for evaluating the risks and determining prices of structured credit products. It may well be that the fantastic growth of this new market has been driven largely by the fact that many consumers of credit ratings failed to appreciate the distinctions between the two types of products superficially labeled to be similar.

Historically, the activities of credit rating agencies had been confined to the corporate finance market, where they would asses a security's quality on the basis of estimates of its anticipated cash flows – its likelihood of default and its expected loss in the event of default. For an individual firm, the two primary considerations in credit risk assessment are the volatility of firm cash flows and the level of outstanding debt (Merton (1974)). Importantly, the precise source of risk in the firm's cash flows – be it firm-specific, industry-wide or from the economy at large – is irrelevant for determining the ultimate rating. This is notable for two reasons. First, the source of the risk determines the level of diversification that can be attained in large securitizations which underlie structured instruments. And, second, understanding the nature of the risk is crucial for the purpose of determining an asset's value, as investors will demand large premia for exposure to systematic risks.

Because structured products are essentially derivatives written on large portfolios of credit-sensitive instruments, estimating cash flows and assigning ratings for structured products becomes far more complicated. Aside from assigning a probability of default for the individual securities in the underlying portfolio, the credit rating agencies are forced to take a stand on the *joint* likelihood of multiple defaults. In effect, in order to assign a rating to a structured product (e.g. CDO tranche) the agency has to be able to model the entire portfolio loss distribution, or equivalently, the dependence structure of defaults (e.g. correlation).[3] This has forced rating agencies to take views on the amount of systematic (or undiversifiable) risks inherent in various underlying assets. By stepping into this new market, the credit rating agencies were required to

[3] Copula-based methods are a popular approach to modeling default dependence (Schonbucher (2003)). The Vasicek model – commonly-used as a market quoting convention – yields a Gaussian copula, in which default dependence is uniquely described by *correlation*.

develop new capabilities related to the estimation of default correlation, with which they had little prior experience.

The systematic risk exposures of structured finance products give rise to a second, and potentially more problematic, distinction vis-à-vis single-name credit securities. Because the payoffs of structured finance securities typically reference a large portfolio of underlying claims, much of the idiosyncratic risk has been diversified away, leaving mainly systematic risk. Since a greater fraction of the total risk of structured securities is derived from systematic risk exposure, these securities should command a yield premium relative to alternatives with equivalent total risk. This is likely to make such securities attractive to investors "searching for yield." Investors who do not fully appreciate or internalize the difference in the nature of the risks between the two types of AAA-rated products are likely to overpay for synthetically manufactured AAA-rated securities, which may, at least partly, explain the extremely rapid growth of this market.

This paper will examine the extreme growth of structured finance market in the context of the above distinctions. We will also consider whether the market structure that evolved over the past century for single-name corporate bonds is likely to be adequate for structured securities.

Credit Rating Agencies

For the better part of the past century, rating agencies have published opinions on the creditworthiness of firms and securities. To reach their opinions, which are labeled as ratings, they gather and analyze a wide range of financial, industry, and economic information. This information is synthesized and published to offer independent, credible assessments of creditworthiness. In exchange for this service, the agencies typically earn fees paid by issuers interested in having their securities rated. By helping resolve information asymmetries between the seller and potential buyers, ratings can deepen a security's pool of potential investors and thereby improve pricing and liquidity. Moreover, to the extent that a firm faces the risk of being downgraded, ratings also have the potential to discipline manager behavior and limit their ability to shift risks.

During the past several decades, credit ratings have come to play a central role in the financial system and are relied upon by a number of key market participants. Many investors use ratings as a starting point for classifying and understanding securities and their associated risks. Lenders or counterparties interested in imposing greater responsibility on a firm may indeed require it to be rated. Buy-side firms such as pension funds, mutual funds, and insurance companies will often use ratings in aiding their investment process. They might use ratings to corroborate, enhance, or even substitute for in-house analysis. They may investigate the reasoning behind a rating to understand and anticipate the likelihood of future actions that might impact prices. And many institutions use ratings in their investment policies – for example, as part of asset allocation rules that require a specific fraction of capital to be invested in securities with certain ratings. Finally, they are likely to use ratings to identify acceptable counterparties in certain transactions and to ensure compliance with regulations.

Regulators also rely on ratings. For example, Rule 2a-7 of the Investment Company Act specifies that money market funds can only purchase commercial paper if it is of sufficiently high rating. Many countries use credit ratings to restrict the kinds of fixed-income investments public and private pension funds can make.[4] Basel II has been modified to allow certain banks to use external credit ratings to determine appropriate capital requirements.[5] Finally, private contracts often rely on ratings to provide external signals that trigger certain actions. For example, a firm might give counterparties the right to demand cash collateral in the event of a downgrade and lenders the right to demand repayment of loans.[6]

---

[4] See Survey of Investment Regulations of Pension Funds, OECD July 2007.

[5] See Benford and Nier (2007) for discussion of use of external credit ratings in Basel II.

[6] Bhanota and Mello (2006) analyze the incentive effects of ratings triggers in corporate debt.

In spite of their wide use, credit ratings have several shortcomings. First, credit ratings only contain information about cash flow risk – default likelihood and expected recovery – but nothing about systematic risk exposure. This information, while important, is insufficient for pricing, since it neglects how payoffs covary with economic states. As a result, ratings create a wide window within which securities of a given rating can have markedly different yields. And although this has historically not been a major issue in the single-name market, with most corporate bonds having similar systematic risk exposures[7], as we discussed above and will demonstrate below, structured finance creates large distortions in securities' systematic risk exposures relative to similarly-rated single corporate bonds.

A second criticism of credit ratings that has been around almost as long as ratings themselves is that they are something of a "sideshow" when it comes to pricing and perhaps even risk assessment itself. In fact, even at the time that the US Securities and Exchange Commission granted the rating agencies their special designation of Nationally Recognized Statistical Rating Organization (NRSRO) there was concern over the dearth of evidence that they were effective.[8] Since then, large literatures emerged that examine the extent to which credit ratings contain information not already incorporated in stock and bond prices.[9] As a whole, the evidence paints at best a mixed picture of the value of ratings. For example, spreads on credit default swaps react modestly to ratings activity[10] and credit ratings do not seem to include much information that is helpful in predicting default beyond that which is already contained in yield spreads.

---

[7] See, for example, Fama and French (1993).

[8] Katz (1974) was among the first studies to examine the efficiency of bond prices with respect to ratings information.

[9] Studies that investigate the stock market reaction to ratings changes include Holthausen, and Leftwich (1986), Hand et al. (1992), Ederington and Goh (1998), Dichev and Piotroski (2001), and Jorion and Zhang (2006). Studies using examining bond returns associated with ratings changes include Hand et al. (1992), Wansley et al. (1992), and Hite and Warga (1997).

[10] Hull et al. (2004) find that credit default swap spreads do not react to ratings downgrades but do react when firm's are placed on review for downgrade.

Indeed, the best models of default likelihood are option-based models that are driven by stock prices (Bharath and Shumway (2007)).

A third concern emerged during investigations of the collapse of Enron – the lack of transparency in the ratings process. Many regulators and issuers argued that increased transparency would allow for greater public scrutiny of the activities of rated companies and reduce the potential for conflicts of interest and the use of fraudulent accounting practices. This resulted in an environment in which issuers designing complex securities had an excellent idea of what models and methods the rating agencies would use to determine their ratings. And, of course, this placed a heavy burden on the rating agencies to avoid any mistakes in their process that could be identified and exploited by savvy issuers.

A final concern is the possibility of conflicts of interest that may arise because the issuer rather than the investor pays for the rating. Mason and Rosner (2007) argue that the process of creating structured finance products requires rating agencies to essentially "become part of the underwriting team," creating large conflicts and producing little transparency for the end investor. The Committee on the Global Financial System from the Bank of International Settlements (BIS) investigated this concern in 2005, and concluded that this potential problem was no more severe for structured finance products than for single-name credit products, arguing that reputation was a strong force against bad behavior in both markets. "In fact, there appear to be no fundamental differences in the rating processes for structured finance products and traditional bonds. The potential conflicts of interest arising in structured finance are thus unlikely to be materially different from those in the traditional segments of the agencies' business" (page 26).

With the recent emergence of structured finance in international capital markets, the role and importance of credit rating agencies expanded dramatically. Because structured finance typically entails securitizations of credit securities (bonds, loans, mortgages), which are themselves often rated, the credit rating agencies were naturally called upon to rate the derivative securities. Indeed, issuers were eager to have their new products rated on the same scale as bonds so that investors subject to ratings-based constraints by their investment mandates would be able to purchase the securities. By having the securities rated, the issuers created an illusion of comparability to existing single-name security which naturally increased the pool of potential buyers for what otherwise would have been perceived as very complex securities.

Rating structured securities became a big part of rating agencies' business, and now accounts as much revenue as their traditional rating of single-name securities.[11] According to Moody's Corporation, 32% of their 2006 revenues came from rating corporate bonds, while 44% of their revenues came from rating structured finance products. This fact is particularly impressive in the context of the relative amount of time it took to develop each revenue source: rating agencies have essentially a decade's experience rating structured finance products whereas they built their single-name rating business over the past century.

Manufacturing AAA-Rated Securities

Manufacturing securities of a given credit rating requires tailoring their cash-flow risk – as measured by the likelihood of default and loss given default – to satisfy the guidelines set forth by the credit rating agencies. Structured finance allows originators to accomplish this goal by means of a two-step procedure comprised of *pooling* and *tranching.* In the first step, a large collection of credit sensitive assets (bonds, mortgages, loans, etc.) is assembled in a portfolio, typically referred to as a special purpose vehicle (SPV).[12] This step alone does not modify the credit rating of the portfolio, as the expected portfolio loss is equal to the mean expected loss on the underlying securities. In other words, the portfolio's credit rating is given by the average rating of the securities in the underlying collateral pool. This situation applies to *pass-through securitizations*, which simply provide their investors with identical, fractional claims to the

[11] For a discussion of the role of structured finance in rating agency revenues see, " The Ratings Charade" by Richard Tomlinson and David Evans. Bloomberg Markets July 2007.

[12] The SPV is generally separate from the originator's balance sheet, so as to isolate the credit risk of its liabilities from the balance sheet of the originator.

payoff on the underlying portfolio. An important example of such structures is provided by *asset-backed securities*, of which residential mortgage-backed securities (RMBS) are a common implementation. By contrast, in order to manufacture securities with customized cash flow risks, structured finance issues a prioritized capital structure of claims, known as *tranches*, against the underlying collateral pool. Because the tranches are prioritized in how they absorb losses from the underlying portfolio, their cash flows – and consequently ratings – can differ materially from the underlying collateral pool. For example, senior tranches only absorb losses after the junior claims have been exhausted, allowing them to obtain credit ratings in excess of the average rating on the collateral pool ("credit enhancement"). The degree of protection offered by the junior claims, or "overcollateralization," plays a crucial role in determining a tranche's credit rating, through its impact on the minimum portfolio loss necessary to impair a senior claim.

The process of pooling and tranching, common to all structured securities, can be illustrated with a simple two-asset example, which also highlights the crucial role of the default dependence structure – a dimension entirely absent from single-name credit rating estimates. Consider two identical securities (e.g. bonds) both of which have a probability of default $p_{D,}$ and pay $0 conditional on default and $1 otherwise. Suppose we pool these securities in a portfolio, such that the total notional value of the underlying fund is $2, and then issue $2 worth of claims (called "tranches") against this fund. A first, "junior" tranche can be written such that it bears the first $1 of losses to the portfolio, paying $1 if both securities avoid default and zero if either security defaults. The second, "senior" claim, which only bears losses if the capital of the junior tranche is exhausted, only defaults if *both* securities default. Intuitively, in order to compute the expected cash flows, or default probabilities, for either of the tranches we will need to know the likelihood of observing both bonds defaulting simultaneously. In this example, the default dependence structure can be succinctly described either by means of a single parameter – either the joint probability of default, or the default correlation.[13]

What makes this simple structure interesting is that if the two underlying securities are imperfectly correlated, the senior tranche will look identical to the original securities except it will have a lower probability of default. For example, if the two securities have a 10% default probability and defaults are uncorrelated, the senior tranche will have only a 1% chance of default. A potentially attractive feature of this basic procedure is that it allows highly risky securities to be repackaged and sold to investors seeking only safe investments.

The above example highlights the two key inputs that allow the securitization machinery to operate. First, multiple securities are required to create sufficiently safe senior assets. In the above example, the senior tranche has a default rate of 1%. For investors unwilling to assume this level of risk, a third security could be added to the structure so that the senior-most $1 of capital only defaults when all three securities default. If each has a 10% probability of default and defaults are independent, this will occur only 0.1% of the time. The second key dimension is the default dependence structure, or equivalently, the (pair-wise) default correlations. The lower the default correlations for the underlying assets, the more improbable it is that all assets default simultaneously and therefore the safer the senior-most claim can be made.

In practice, the creation of structured credit products follows the above template relatively closely. For example, a typical collateralized loan obligation (CLO) will have as its underlying a portfolio of 100-200 different bank loans purchased at issue over a 6 month period. Each of these loans might be rated BBB, suggesting that each loan has about a 2% to 5% chance of defaulting during a five-year period using historical frequencies. Since bank loans are the most senior claims that corporations can issue, they have historically had relatively high recovery values in the instances where they default, averaging around 60% of their notional value (Altman

---

[13] If we assume that both securities are identical and denote the probability of observing both claims default simultaneously by pDD, the *default correlation* parameter can be computed as $(p_{DD}-p_D^2)/(p_D*(1-p_D))$.

and Swanson (2007)). Next, a prioritized capital structure is issued against the cash flows generated by the underlying portfolio. The claims, or tranches, issued against the portfolio are defined according to how the portfolio losses are to be distributed. The claim exposed to the first losses on the portfolio is typically called the "equity tranche." The claims exposed to the subsequent portfolio losses are typically called "mezzanine tranches." Finally, the "senior tranches" are exposed to portfolio losses only after the equity and mezzanine tranches have been completely exhausted. An illustration of this hypothetical Collateralized Loan Obligation (CLO) is presented in Figure 1.

Figure 1
Hypothetical CDO Structure

Reference Entity	Notional Value	Loss Exposure	Tranche Name (Rating)	Notional Value
Portfolio of Individual Credit Securities 150 Loans; Avg Rating of BBB	100	20%-100%	Senior (AAA)	80
		5%-20%	Mezzanine (BBB)	15
		0%-5%	Equity (NR)	5
Assets	100		Liabilities	100

- If portfolio losses are less than $5, then the Mezzanine and Senior tranche receive promised payments, while the Equity tranche suffers the realized loss.
- If portfolio losses are between $5 and $20, then the Senior tranche receives promised payments, the Mezzanine tranche suffers the realized loss less $5, and the Equity tranche is exhausted.
- If portfolio losses exceed $20, then the Senior tranche experiences the realized loss less $20, and both the Mezzanine and Equity tranches are exhausted.

Rating Structured Finance Products

As the above examples highlight, assigning ratings to CDO tranches requires the credit rating agency be able to characterize the entire joint-distribution of payoffs for the underlying collateral pool. This distribution depends crucially not only on the estimates of default probabilities and loss values for the underlying claims (i.e. their credit ratings), but also on the default dependence structure. Only with this information, can the credit rating agency determine the probability that the loss on the underlying portfolio will exceed the level of protection offered by the claims junior to the one being rated, precipitating its default.

The requirement that the credit rating identify the true underlying default dependence structure demands a substantially increased level of rigor and consistency in the ratings of the underlying securities. For example, although the rating of the individual security will not depend on the source of the risk in its underlying cash flows – firm-specific, industry-wide or from the economy at large – the partitioning of these sources of risk is crucial for arriving at the correct estimates of default dependencies (correlations). Exposures to factors that are idiosyncratic to individual firms will have no effect on driving the joint likelihood of defaults, whereas exposures to industry- and economy-wide factors will. Consequently, even if the credit ratings agencies correctly assign ratings to individual securities, there is substantial scope for error in the ratings of the derivative, structured products, if the ratings process fails to partition the exposure to the undiversifiable factors driving default correlation. Finally, the sequential application of capital structures in structured products (e.g. CDOs of CDOs), can accentuate errors in assignments of default probabilities at the level of individual securities.[14] Even minute errors at the level of the underlying securities, which would be insufficient to alter the security's rating, can become extreme with multiple rounds of structuring, such that ratings on structured securities can be significantly off.

Characterizing the entire default dependence structure for the securities in the underlying portfolio is extremely daunting. The underlying collateral pools are oftentimes comprised of more than one hundred securities, forcing the credit rating agency to take a stand on nearly 5000 pairwise default correlations (N = 100).[15] To sidestep this curse of dimensionality, rating agencies are generally forced to assume that the underlying securities are identical within certain

[14] CDOs whose underlying asset pool consists of tranches of other CDOs appear to have been especially common in the structured finance products involving mortgages.

[15] The number of pairwise correlation is given by *Nx(N-1)/2*, where *N* is the number of securities in the underlying portfolio.

groupings (e.g. industries), such that only estimates of within and across industry default correlations are necessary. Unfortunately, due to the rarity of defaults, many of these assumptions pertain to features of the data that are either unobserved and/or poorly measured.

Intuitively, default correlations are driven by common exposures of the underlying securities to systematic (non-diversifiable) risk factors. For example, although the performance of any two mortgages originated in a given zip code depends on myriad considerations which are specific to each borrower's situation, the mortgages will also be crucially affected by common variation in house price appreciation, or the state of the local economy.

Any errors in estimating the exposures to the common factors, feed through to the estimates of default correlation, ultimately biasing the ratings assigned to the CDO tranches. Depending on the estimates of the default correlation, the loss distribution for the underlying portfolio can take on a wide range of shapes, affecting the expected loss on various claims in the capital structure. To illustrate the effects of correlation on the loss distribution of an underlying portfolio, let us return to our simple CLO example. We assume that that underlying portfolio is comprised 100 identical loans, each with a default probability of p = 0.05, and a (deterministic) recovery value of 0.60 cents on the dollar conditional on default. In Figure 3, we present the impact of default correlation on the loss distribution. As the figure makes clear, increasing the default correlation increases the possibility of extreme losses – both very small and very large. When defaults are uncorrelated (rho = 0.00), the distribution takes on a symmetric, bell shape reflecting the intuitive fact that extreme outcomes, in which very large or very small number of firms default are rare. As correlation increases, so does the likelihood of extreme outcomes. In the limit, as default correlation becomes perfect (rho = 1.00), either all firms are expected to survive or default, resulting in a bimodal loss distribution that has peaks around extreme losses.

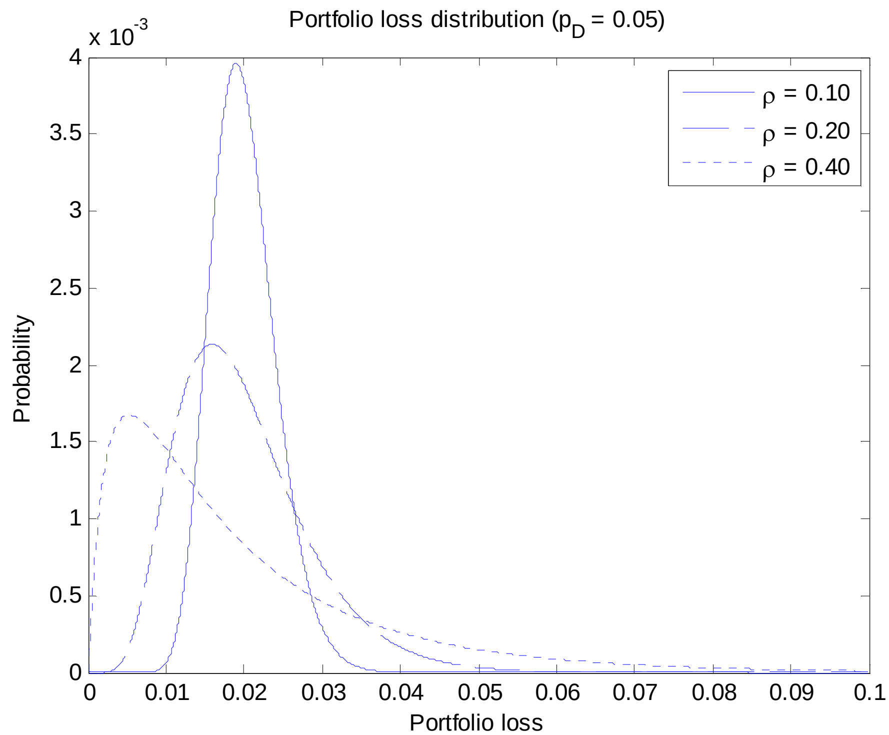

To examine how this impacts the various tranches in the capital structure of a structured product, we focus on the three tranches in our example – an equity tranche, a mezzanine tranche and a senior tranche. The CLO is designed such that the equity tranche absorbs the first 5% of the losses on the underlying portfolio, the mezzanine tranche – the following 15% of the losses, and finally the senior tranche is only affected once the cumulative loss on the underlying portfolio exceeds 20%. Under the parameters used to characterize the loans in the underlying portfolio, the expected portfolio loss is equal to 2% (= 0.05 * (1-0.60)) of the outstanding notional. If the portfolio loss was always guaranteed to be at its expected value, only the equity tranche would sustain a loss. Of course, in reality actual losses can vary substantially from their *ex ante* expectation, causing the more senior claims to be affected. In Figure 4, we explore the impact of the correlation in the underlying asset returns on the expected tranche losses.[16] The figure illustrates that increases in correlation shift the distribution of cash-flow risk across the

---

[16] Default correlations are monotonically related to the asset return correlations. The precise relationship can be derived in closed-form under the Vasicek model (1991).

tranches, effectively making the junior claims safer at the expense of the senior claims. All the while, the total loss sustained by the CLO tranches is determined by the features of the underlying loans (default probability, recovery rates conditional on default) and always sums to 2%, independent of the correlation value.

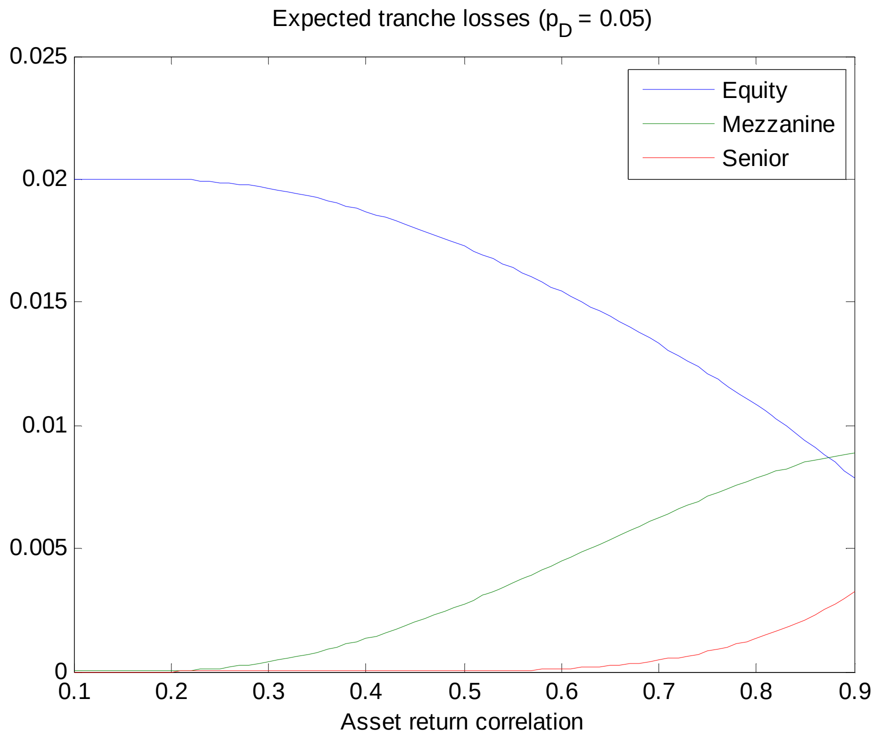

The results presented thus far highlight the importance of proper estimation of correlation in getting the relative pricing of the various tranches of a structured product right. Other things equal, an investor that under-appreciates – or relies on ratings that have underestimated – the degree of correlation in the underlying asset pool will find equity tranches relatively expensive and senior tranches relatively inexpensive.

Given the central role of correlation in allocating expected losses across the various claims in the capital structure, constructing reliable estimates of default correlations is crucial for assigning credit ratings and determining prices. Unfortunately, arriving at sensible default correlation estimates is extremely difficult. As a result, rating agencies and Wall Street practitioners have developed several methods for estimating these values. For example, Moody's Investors Service (2004) outlines three main approaches to estimating default correlations. The first approach is to use historical data on actual defaults. This approach, while most directly relevant to the variable of interest, is difficult to implement because there are not enough actual defaults in the historical database to accurately extract an estimate of how these events correlate with each other. A second possibility is to use market prices and a pricing model to imply the capital market's estimate of the average default correlation. This method is widely used by Wall Street practitioners who are largely in the business of market making. They need to have a sense of how the market is pricing securities, but are not in the business of determining whether the market prices are accurately reflecting all available information. Much in the way that the Black-Scholes option pricing model and option prices can be used to determine an "implied volatility," a credit pricing model and tranche prices can be used to determine implied correlations once expected losses have been fixed. This approach is also used to some degree by rating agencies. For example, it is relied upon by Moody's KMV, which pioneered the use of extracting the capital market's views about various parameters from market prices and a structural pricing model. The same model that is used by KMV to construct its expected default frequency (EDF) can be used to estimate asset return correlations, which can subsequently be translated into default correlation estimates (GCorr).

A final approach for estimating default correlation is to use credit rating transitions as an event dataset from which to estimate the correlation of actual default. This method has been adopted as the preferred approach by at least some of the rating agencies. For example, by examining rating transitions, Moody's arrived at a value of 0.15 for intra-industry default correlation and 0.03 for cross-industry default correlation. Although virtually identical intra-industry default correlations are implied by market prices using the approach described, the implied inter-industry default correlations are considerably higher, on the order 0.13, according to Moody's (2004). Relatedly, default correlations can also be estimated from the co-movements of credit default swap spreads. Unfortunately, the CDS market is relatively new, such that data limitations – both in terms of the cross-section of covered securities and in the available time series – severely restrict the ability to apply this method.

Overall, due to the highly non-linear impact of correlation on the valuation of the various tranches, seemingly small discrepancies in default correlation estimates can have a significant impact on assessing tranche risks, and consequently, assigning their ratings. Without the benefit of prior experience in tackling the issue of estimating default correlations, it is highly questionable that credit ratings agencies have been able to successfully resolve the wide range of underlying problems. In turn, one must be skeptical in appraising claims made by the rating agencies, that corporate and structured securities with identical credit ratings are indeed likely to be comparable in terms of their *ex post* default experience.

Economic Catastrophe Bonds

An additional puzzle concerning the reliance upon ratings by investors is that, in the context of structured finance, they may not be that helpful for valuation. Moreover, to the extent that investors rely upon ratings for valuation, highly-rated structured finance products are likely to be viewed as attractive to cautious investors seeking safe securities, but who do not fully appreciate or internalize the nature of the risks they are bearing.

The central insight of asset pricing theory is that in order to determine an asset's price one has to know both its expected payoff and how that payoff covaries with risks investors care about. The higher the expected payoff, the higher the asset's price, and conversely, the higher the payoff covariance, the less insurance the asset provides against systematically bad states and the lower the asset's price. As discussed above, credit ratings only assess a security's expected payoff (i.e. likelihood of default and anticipated recovery rate) – they offer no information regarding the degree to which a security pays off in different states.

Because credit ratings only reflect expected payoffs, there can be a wide range of yield spreads (yield in excess of a maturity or duration matched Treasury security) that a security with a given credit rating can command depending on its systematic risk exposure. The minimum of this range is set by the expected loss rate and consists of securities that have identical likelihood of default across economic states.[17] An example of a security whose yield spread is close to its expected loss rate is a traditional catastrophe bond that defaults in the event of a natural disaster such as a hurricane or earthquake. Since a single natural disaster cannot have a material impact on the world economy[18] its risks are negligible in the context of a well-diversified portfolio. As a consequence, a traditional catastrophe bond earns yield solely as compensation for expected losses and warrants no additional yield premium to compensate for risks that, while considerable, are imminently diversifiable.

At the other end of the range of yield spreads for securities of a given rating are securities that default primarily in bad economic states. The maximum yield is set by a security that confines defaults entirely in the worst economic states. If the stock market can be thought of as providing an ordering of economic states (i.e. if the S&P is at 800, the economy is in worse condition than if the S&P is at 900), then the security with maximal exposure to systematic risk is a digital call option on the market. A digital call option is a security that pays off $1 if the market is above a pre-determined level (called a "strike price") at maturity and $0 otherwise. Because this security "defaults" (i.e. fails to pay) only and always when the market is below the strike price, it represents the security with the greatest possible exposure to systematic risk. By setting the strike price appropriately, the digital call can have a loss rate equal to securities with a given credit rating. But because it concentrates default in the worst economic states, its price will be heavily discounted to deliver a yield that is maximal for that rating.

The process of pooling and tranching effectively creates securities whose risk profiles resemble those of a digital call on the market index – or, more precisely a call spread on the market index. To see this, notice that pooling allows for diversification, such that the

[17] In theory, a bond that is more likely to pay off in expensive than in inexpensive states can have a yield that is below its expected loss rate.

[18] Even Hurricane Katrina, whose $100 Billion estimated cost makes it the most costly natural disaster in US history, is small in the context of an economy whose GDP alone is over 100 times as large, and much less in the context of the global economy.

contribution to portfolio loss attributable to idiosyncratic (firm-specific) asset return realizations declines. In the limit, losses are driven entirely by the underlying systematic risk exposure. As a result, a tranche written against a large and highly diversified pool of underlying assets is essentially the same as a derivative written against the market index. Because the tranche begins losing value when losses on the underlying pool have exceeded the upper attachment point, and is fully exhausted when losses exceed the lower attachment point, the tranche payoffs resemble those of a call spread on the market.

In effect, structured finance has enabled investors in the senior claims to write insurance against large declines in the market. To the extent that investors do not fully appreciate the nature of the insurance they are writing, they are likely to be earning yield that is attractive relative to that of similarly rated securities (i.e. securities with similar likelihood of default), but well below the return they can earn from simply writing such insurance directly in the index options market. Coval, Jurek, and Stafford (2007) provide evidence in support of this conjecture, showing that senior tranches in CDOs do not offer their investors nearly enough yield spread to compensate them for the actual systematic risks they are being asked to bear.

The fact that corporate bonds and structured finance securities carry risks that can, both in principle and in fact, be so different from a pricing standpoint, casts significant doubt on whether corporate bonds and structured finance securities can really be considered comparable, despite their being rated on a similar scale, and indicates that separate regulation may be necessary.

Implications and Conclusions

The big question is how the market for highly rated structured finance products will evolve in the coming years? Many practitioners believe that the current credit crunch will work itself out quickly as markets tend to do, and the market for structured credit will eventually return to its earlier levels. This view faces the challenge that the market for structured credit appears to have serious structural problems that may be difficult to overcome in the short-run.

The immediate challenge is that the buyers of highly rated structured finance products are not buying. The CDO and CLO structures work when the senior tranches can be sold at low spreads (or high prices), but the traditional buyers of these claims are not currently interested at the historical terms. Interestingly, AAA-rated corporate bonds are currently being issued at a rapid pace, suggesting that the concern over AAA-rated securities is confined to those manufactured through structured finance. As we saw earlier, the value of these claims are highly sensitive to the assumptions of (1) default probability and recovery value, (2) correlation of defaults, and (3) the relation between payoffs and the economic states that investors care about most. The buyers of highly rated structured finance products are concerned that essentially all of the assumptions that they have been relying upon have been systematically biased against them, and are now reluctant to invest in securities that they do not fully understand, no longer willing to outsource so much of their credit analysis to the rating agencies. The key question is whether they will again.

The existing market structure relies heavily on rating agencies, who were historically in the business of providing estimates of default probability and recovery values for single-name credit securities using historical data. Estimates of default correlations are not required for single-name credit risk. And while the rating agencies were never in the business of describing how payoffs varied across economic states, markets seem to have developed workable mappings of their own. These mappings are essentially approximations to a proper asset pricing exercise that would involve an integrated analysis of the distribution of expected payoffs across economic states and state prices. By relying on the rating agencies for proxies of expected payoffs for credit securities with no information on how they relate to various economic states, an approximate pricing approach must be used, which is likely to produce systematic errors. So long as the costs of these errors were small enough, investors were willing to rely on rating agencies to avoid having to conduct the fully integrated pricing exercise. In other words, the market for single-name credit risk was reasonably efficient if the proper pricing exercise was sufficiently costly to perform.

Importantly, structured credit securities represent a distinct investment product from single-name credit securities. Therefore, it is not clear that the market structure for single-name credit products can simply be adopted by this new product market. As made clear from the earlier analysis, the risk characteristics of senior tranches of CDOs and CLOs are highly dependant on estimates of default correlations to determine expected payoffs, and again, for proper pricing one needs to know how expected payoffs are distributed conditional on economic state. The fully integrated proper pricing exercise is not materially harder for these securities than it is for single-name credit securities. However, we have seen that the costs associated with these pricing errors are considerably higher, such that we may be at the tipping point where the recent investors in these securities are better off either employing a more integrated pricing exercise themselves, relying less on the rating agencies than they did before; or by avoiding these securities altogether.

References

Altman, Edward and Jeffrey Swanson, 2007, "The Investment Performance and Market Size of Defaulted Bonds and Bank Loans," NYU Stern working paper.

Benford, James and Erlend Nier, 2007, "Monitoring cyclicality of Basel II capital requirements," Bank of England, Financial Stability Paper No. 3.

Bhanota, Karan, and Antonio S. Mello, 2006, "Should corporate debt include a rating trigger?" *Journal of Financial Economics* 79, pp. 69-98.

Bharath, Sreedhar and Tyler Shumway, 2007, "Forecasting Default with the Merton Distance to Default Model," *Review of Financial Studies*.

Coval, Joshua, Jakub Jurek, and Erik Stafford, 2007, "Economic Catastrophe Bonds," Harvard Business School working paper.

Dichev, Ilia D. and Joseph D. Piotroski, 2001, “The Long-Run Stock Returns Following Bond Rating Changes,” *Journal of Finance* 56, pp. 173-203.

Ederington, Louis H., and Jeremy Goh, 1998, “Bond Rating Agencies and Stock Analysts: Who Knows What When?” *The Journal of Finance and Quantitative Analysis* 33, 569-85.

Fama, Eugene F. & French, Kenneth R., 1993. “Common risk factors in the returns on stocks and bonds,” *Journal of Financial Economics* 33, pp. 3-56.

Fitch Ratings, 2007, “Inside the Ratings: What Credit Ratings Mean.”

Hand, John R., Robert W. Holthausen and Richard W. Leftwich, 1992, “The Effect of Bond Rating Agency Announcements on Bond and Stock Prices,” *Journal of Finance* 47, pp. 733-52.

Hite, Gailen and Arthur Warga, 1997, “The Effect of Bond-Rating Changes on Bond Price Performance,” *Financial Analysts Journal* 53, 35-51.

Holthausen, Robert W. and Richard W. Leftwich, 1986, “The Effect of Bond Rating Changes on Common Stock Prices,” *Journal of Financial Economics* 17, pp. 57-89.

Hull, John, Mirela Predescu, and Alan White, 2004, “The Relationship Between Credit Default Swap Spreads, Bond Yields, and Credit Rating Announcements," *Journal of Banking and Finance* 28, pp. 2789-2811.

Jorion, Philippe, and Gaiyan Zhang, 2007, “Good and bad credit contagion: Evidence from credit default swaps,” *Journal of Financial Economics* 84, pp. 860-883.

Katz, Steven, 1974, “The Price Adjustment Process of Bonds to Rating Reclassifications: A Test of Bond Market Efficiency,” *Journal of Finance* 29, pp. 551-59.

Mason, Joseph and Joshua Rosner, 2007, “Where Did the Risk Go? How Misapplied Bond Ratings Cause Mortgage Backed Securities and Collateralized Debt Obligation Market Disruptions,” Drexel University working paper.

Merton, Robert C., 1974, "On the Pricing of Corporate Debt: The Risk Structure of Interest Rates," *Journal of Finance* 29, pp. 449-470.

Moody's Investors Service, 2004, "Moody's Revisits Its Assumptions Regarding Corporate Default (and Asset) Correlations for CDOs."

Organisation for Economic Cooperation and Development, 2007, "Survey of Investment Regulations of Pension Funds."

Schonbucher, Pilipp J., 2003, "Factor Models: Portfolio Credit Risks When Defaults Are Correlated," *Journal of Risk Finance* 3, pp. 45–56.

Vasicek, Oldrich, 1991, "Limiting Loan Loss Distribution," Working Paper, KMV Corporation.

Wansley, James L., John Glascock, and Terrence M. Clauretie, 1992, "Institutional Bond Pricing and Information Arrival: The Case of Bond Rating Changes," *Journal of Business Finance and Accounting* 19, pp. 733-750.