Economic Catastrophe Bonds

Joshua D. Coval, Jakub W. Jurek, and Erik Stafford*

March 2007

Abstract

The central insight of asset pricing is that value depends on both expected payoffs and how payoffs covary with risks investors care about. In credit markets, many investors focus exclusively on measures of expected payoffs (e.g., credit ratings, default likelihoods) without considering the state of the economy in which default is likely to occur. Such investors are likely to be attracted to securities whose payoffs resemble those of economic catastrophe bonds – bonds that default only under severe economic conditions. This paper argues that almost all securities created by structured finance share this feature. We demonstrate that although the risks of these securities closely resemble those of a more transparent economic catastrophe bond – a portfolio that is long a Treasury bond and short deeply out-of-the-money S&P put options – their prices do not. Investors in structured finance products receive far less compensation for bearing economic catastrophe risk than those that write out-of-the-money index put options. We argue that this difference arises from the willingness of rating agencies to certify the former as "safe" and from the large supply of investors who view them as such.

*Coval, Jurek, and Stafford are at Harvard University. We thank John Campbell, Bob Merton, and André Perold for valuable comments and discussions, and seminar participants at Boston College and the Harvard PhD finance seminar. We are especially grateful to Eli Cohen and Marco Naldi at Lehman Brothers and Max Risman and Jeff Larson at Sowood Capital for providing data and insights on credit markets.

Economic catastrophe bonds fail to deliver their promised payments in the "worst" economic states, precisely when a dollar is most valuable. Like most catastrophe bonds, they can easily be constructed to have a small probability of default, or equivalently a high expected payoff relative to face value. However, unlike bonds that insure non-economic or actuarial risks, economic catastrophe bonds should have large price discounts to account for the extreme variation of payoffs across states with different levels of aggregate wealth. This simple, but seemingly abstract security becomes meaningful when one considers recent trends in the development of structured financial products. The increasingly popular act of pooling economic assets into large portfolios and tranching them into sequential cash flow claims effectively creates economic catastrophe bonds. For example, a so-called "safe" tranche of a collateralized default obligation (CDO) is AAA-rated based on its low default probability, but fails to deliver its promised payoff in an economic environment in which 15% to 20% of large US firms have defaulted on their debt, which is almost surely a state where aggregate wealth is low. The analysis presented in this paper shows how recent innovation in financial products tends to concentrate risks in a manner that is particularly important for pricing, but which can be obscured with common methods of credit risk and price analysis.

Structured finance instruments have played a key role in the recent growth of financial markets, with $3.1 trillion in issuance in 2005 alone and a total outstanding, notional value of $8.1 trillion at 2005 year-end. These instruments are typically created in a process known as securitization, in which assets are first combined into a collateral pool, and then claims with varying degrees of seniority, known as tranches, are issued against this pool. One of the motivations for securitization is its ability to create derivative claims with credit ratings superior to the average credit rating of the assets in the underlying pool. This allows for an expansion of the investment opportunity set by creating payoffs that were previously unavailable or available naturally in short supply.

As a result of the prioritization of cash flows (liabilities) stemming from the underlying collateral pool, a senior claim only suffers losses after the principal of all of the subordinate tranches has been exhausted. While this prioritization scheme effectively ensures that the senior tranche has a high expected payoff, and is likely to receive a commensurately high credit rating, it also confines the tranche losses to systematically bad states. 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 (e.g. credit rating), 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 in systematically bad states and the lower the asset's price. Interestingly, while asset pricing theory predicts that the AAA-rated tranche should trade at a relatively low price, commanding a large risk premium, we show in this paper that this prediction does not find empirical support. This finding not only suggests the presence of a quantitatively large mispricing, but also highlights the need for a careful investigation of the structure of the credit market, including the role of credit ratings in the structured finance domain.

Credit agencies have long played a central role in fixed income markets, where they categorize investments in terms of their likelihood of meeting coupon and principal repayment obligations.

Although these ratings do not include any form of adjustment for how default is expected to relate to other economic risks, they are heavily relied upon by investors, issuers, and monitors in valuing and assessing the risks of fixed income investments. In the case of single-name issuers, neglecting the covariance of the bond's payoff with other economic risks turns out to be a relatively benign shortcut, since the majority of the cross-sectional variation in bond yields is attributable to differences in payoff expectations, rather than risk premia. In other words, a naïve investor who prices bonds solely on the basis of their credit rating would not be making a large mistake. However, because pooling and tranching can have a dramatic impact on the risk characteristics of the resulting claims, such shortcuts are generally no longer warranted for structured finance instruments.

To investigate the recent behavior of the structured credit market we develop a simple framework for understanding how tranching schemes commonly applied to portfolios of economic assets affect risk and pricing. In particular, we show that if investors were to naïvely price securities using rules of thumb that heavily emphasized expected payoffs over appropriate required rates of return, then the tranching of economic assets, or the creation of economic catastrophe bonds naturally emerges as an optimal means of exploiting these investors. Our empirical analysis suggests that investors that purchase senior CDO tranches receive yield spreads that are several factors too low to appropriately compensate them for the highly systematic nature of the risks they are bearing. In particular, we estimate that an investor that purchases the AAA-rated CDO tranche bears risks that are highly similar to those of a 50% out-of-the-money five-year put spread on the S&P 500 index. However, on average the put spread offers nearly three times as much compensation for bearing these risks.

The remainder of the paper is organized as follows. Section 1 develops a simple framework for understanding how tranching schemes commonly applied to portfolios of economic assets affect risk and pricing. Section 2 describes the data. Section 3 presents a calibration methodology that allows us to compare the risk and return properties of CDO tranches to market index put spreads that have equivalent default risk. Section 4 evaluates the time series properties of actual CDO tranche spreads relative to model predicted spreads. Section 5 discusses the recent evolution of the structured credit market, and Section 6 concludes.

1 The Impact of Tranching on Asset Prices

Assets cannot be priced solely on the basis of their expected payoff. This simple insight underlies the entirety of modern asset pricing, which stipulates that in order to determine the price of an asset one has to know both its expected payoff, and how that payoff covaries with priced states of nature (i.e. the stochastic discount factor). Take, for example, the case of a risky discount bond which pays one dollar T-periods hence, conditional on not defaulting, and zero otherwise. The price of this bond can be obtained from the fundamental law of asset pricing, which states that an asset's price is given by the expectation of the product of its future payoff, CF_T, and the realization of the stochastic discount factor, M_T,

$$P_0 = E[CF_T \cdot M_T] = e^{-r_f \cdot T} \cdot E[CF_T] + Cov[CF_T, M_T] \tag{1}$$

In the case of this risky discount bond, which pays zero conditional on default, the future cash flow is given by,

$$CF_T = (1 - \mathbf{1}_{D,T}) \cdot 1 + \mathbf{1}_{D,T} \cdot 0 \tag{2}$$

where $\mathbf{1}_{D,T}$ is an indicator random variable, which takes on the value of one conditional on the bond being in default at time T, and zero otherwise. If the probability of default at time T is given by $\overline{p}_D$, the bond's price will satisfy,

$$P_0 = e^{-r_f \cdot T} \cdot (1 - \overline{p}_D) - Cov[\mathbf{1}_{D,T}, M_T] \tag{3}$$

The bond price is equal to the the expected future cash flow discounted at the riskless rate, adjusted for the covariation of defaults with priced states of nature. Although the relative magnitude of the two terms is likely to vary across various securities, the rapid growth of credit rating agencies, which specialize in delivering unconditional estimates of default probabilities and losses given default, suggests that practitioners are most interested in the first term. Of course, there are circumstances where this shortcut can lead to significant errors. The pricing formula, (3), reveals that neglecting the risk premium for the covariation of defaults with prices states of nature may lead to severe mispricings. In particular, we argue that the magnitude of the potential mispricing is likely to be largest within structured finance products, where the risk premium is magnified through the pooling and tranching of securities. Paradoxically, the largest recent driver of credit rating agency revenues – structured finance products (e.g. collateralized debt obligations) – are also likely to be the products where estimates of default probabilities are least likely to be sufficient for pricing.

In the next section we provide some intuition for the magnitude of the mispricing that can be created by neglecting the risk premium for covariation of defaults with priced states of nature, and show how pooling and tranching reallocates payoffs across these states. Indeed, if market participants assigned identical prices to all fixed income securities with identical credit ratings, issuers would have an incentive to create and sell securities whose default probability strongly covaries with priced states of nature. We show that tranching arises as an endogenous mechanism for exploiting this naïve, credit-rating-based approach to pricing fixed income securities.

1.1 The Cheapest to Supply Bond

To get a sense of how much the prices of a set of bonds with identical unconditional default probabilities, i.e. credit ratings, can vary, let us consider all possible payoff profiles in the priced state space, $\mathbf{\Omega}$. As before, we will assume that the bond either pays one dollar conditional on not defaulting, and zero otherwise. If we denote the state-contingent probability of default by $p_D(\omega)$

and the probability of observing state ω by $f(\omega)$, this set of securities includes all bonds that satisfy,

$$\overline{p}_D \equiv \int_{\omega \in \mathbf{\Omega}} p_D(\omega) f(\omega) d\omega \tag{4}$$

where $\overline{p}_D$ is the pre-specified, unconditional default probability. In principle, to price these securities we can simply integrate their state contingent payoff expectations against the state prices, $q(\omega)$,[1]

$$P(p_D(\omega), \overline{p}_D) \equiv \int_{\omega \in \mathbf{\Omega}} (1 - p_D(\omega)) q(\omega) d\omega. \tag{5}$$

However, to derive bounds on the prices of the bonds it is useful to re-write the above expression in terms of the stochastic discount factor, $m(\omega)$, which is given by the ratio of the state price and the state probability, $f(\omega)$,

$$\begin{aligned} P(p_D(\omega), \overline{p}_D) &= \int_{\omega \in \mathbf{\Omega}} (1 - p_D(\omega)) \cdot \left(\frac{q(\omega)}{f(\omega)} \right) \cdot f(\omega) d\omega \\ &= \int_{\omega \in \mathbf{\Omega}} (1 - p_D(\omega)) \cdot m(\omega) dF(\omega). \end{aligned} \tag{6}$$

The stochastic discount factor, $m(\omega)$, reflects the marginal utility of consumption in each state and provides a natural means by which states can be ordered from "most expensive" to "least expensive". Once the states ω have been ordered according to their corresponding value of $m(\omega)$ – from highest to lowest – it is immediate that the most expensive asset pays off with certainty on a set of measure, $1 - \overline{p}_D$, containing the most expensive states – as measured by $m(\omega)$ – and zero elsewhere. We denote the set of states in which the most expensive asset delivers a unit payoff by $\underline{\mathbf{\Omega}}$. Conversely, the least expensive asset pays of with certainty on a set of measure $1 - \overline{p}_D$, but containing the least expensive states, and zero elsewhere. Correspondingly, we denote the set of states with a sure, one unit payoff for the cheapest asset by $\overline{\mathbf{\Omega}}$. However, absent an explicit characterization of the priced state, ω, and the state-contingent value of the stochastic discount factor, $m(\omega)$, it is not possible to determine how large the wedge is between the prices of these two identically rated assets.

One natural state space in which to consider the pricing of these 'toy' securities is the state space defined by the realizations of the market return. This state space underlies the Sharpe (1963) - Lintner (1965) capital asset pricing model, and plays a crucial role in many other multi-factor characterizations of priced states. Moreover, because the market factor describes the evolution of wealth of the representative agent, low (high) realizations of the market return identify states with high (low) marginal utility, or equivalently, high (low) values of the stochastic discount factor. Consequently, indexing states by the magnitude of the realized market return also provides the requisite ordering of states in descending order of marginal utility of consumption.

In the market state space, the two securities with the highest and lowest prices, and an uncon-

[1] The state price is equal to the price of the Arrow-Debreu security for state ω, i.e. a security which pays one unit of consumption in state ω and zero otherwise.

ditional default probability of $\overline{p}_D$, correspond to a digital market put and call option, respectively. The strike price of each option is set such that the probability of observing the option expire out of the money is equal to $\overline{p}_D$, and the sets of states for which they yield unit payoffs correspond to the perviously identified $\underline{\mathbf{\Omega}}$ and $\overline{\mathbf{\Omega}}$. To price these options analytically it is convenient to specialize to the assumptions underlying the Black-Scholes (1973) / Merton (1973) option pricing model, and require that the market follow a lognormal diffusion with constant volatility. Under this specification, the evolution of the market is described by the following stochastic differential equation,

$$\frac{dM}{M} = (r_f + \eta)dt + \sigma_m dZ_m. \tag{7}$$

where, r_f is the continuously compounded riskless rate, η it the market risk premium and σ_m is the volatility of instantaneous market returns. Moreover, the prices of Arrow-Debreu securities, and many other derivatives, including digital options, can be obtained in closed form (see Breeden and Litzenberger (1978)). After some simple manipulation it is possible to show that the prices of the digital market put and call option with a default probability of $\overline{p}_D$ are given by,

$$\overline{P}_0 = e^{-r_f T} \cdot \Phi\left[\Phi^{-1}(1-\overline{p}_D) + \frac{\eta}{\sigma_m} \cdot \sqrt{T}\right] \tag{8}$$

$$\underline{P}_0 = e^{-r_f T} \cdot \Phi\left[\Phi^{-1}(1-\overline{p}_D) - \frac{\eta}{\sigma_m} \cdot \sqrt{T}\right] \tag{9}$$

where $\Phi(\cdot)$ denotes the cumulative normal distribution of the standard normal random variable. These expressions indicate that the maximal and minimal prices for a bond with an unconditional default probability of $\overline{p}_D$, will depend on the default probability itself (i.e. expected cash flow), and the T-period market Sharpe ratio. As intuition would suggest, when the market Sharpe ratio is equal to zero (i.e. no risk premium), the prices of the two bonds will be identical, and equal to the price of a discount bond with a constant, idiosyncratic default probability of $\overline{p}_D$ in all market states.

To get a sense of the magnitude of the mispricing that can arise from omitting the risk premium in the computation of the price of a security with a 5-year unconditional default probability of 1%, consider the following calibration. Suppose the (annualized) continuously compounded riskless rate, r_f, is equal to 5%, and that the annualized market Sharpe ratio is 0.33. Under these assumptions the price of a discount bond with a par value of one, whose defaults are purely idiosyncratic, would be equal to $P_0 = 0.7710$. On the other hand, the price of the cheapest security with the identical default probability is given by the price of digital market call, $\underline{P}_0$, and is equal to 0.7351. If market participants naïvely assume that defaults are idiosyncratic and assign equal prices to all securities with an identical credit rating, a clever agent could exploit them by obtaining a rating for the digital market call, and marketing it at the price of other securities with the same rating, while pocketing the 4.66% price differential.

This simple analysis illustrates that securities with identical credit ratings, interpreted as unconditional default probabilities, can trade at significantly different prices. This is not surprising in the context of asset pricing theory, which posits that an asset's price should reflect a premium for the covariation of its payoff with priced states of nature. It also suggests a simple mechanism for exploiting market participants who naïvely assign the same price to all securities with the same credit rating. So long as the price assigned to a security of a given credit rating differs from the price of the cheapest to supply bond, i.e. the digital market call, arbitrageurs have an incentive to sell digital market calls, or other securities with similar payoff profiles. However, the transparency of this ploy, combined with the improbability of being able to obtain a credit rating for a digital market call option, suggests this is not possible. Astoundingly, we show that tranching the cash flows from a portfolio which pools a large number of economic assets (e.g. bonds, credit default swaps, etc.) – a commonly accepted market practice aimed at obtaining credit enhancement – does just this.

1.2 Tranching as a Mechanism for Reallocating Risk

Structured finance activities effectively proceed in two steps. In the first step, a portfolio of similar securities (bonds, loans, credit default swaps, etc.) is *pooled* in a special purpose vehicle. In the second step, the cash flows of this portfolio are redistributed, or *tranched*, across a series of derivatives securities. The absolute seniority observed in re-distributing cash flows among the derivative claims, called tranches, enables some of them to obtain a credit rating higher than the average credit rating of the securities in the reference portfolio. Aside from allowing the issuer to satisfy the demands of clienteles with various tolerances for default risk, tranching also mitigates asymmetric information problems regarding the quality of the underlying securities (DeMarzo (2005)). Unlike DeMarzo though, our focus is not on the agency problems motivating the existence of tranching, but rather on its impact on the systematic risk exposures of the resulting securities, and consequently, on their prices. We show that losses on highly rated tranches are concentrated in states with high state prices (i.e. marginal utility), suggesting that they should trade at significantly higher yield spreads than single-name bonds with identical credit ratings. Surprisingly, this implication turns out not to be supported by the data, which shows that triple-A rated tranches trade at comparable yields to triple-A rated bonds. This suggests a different, and more tantalizing, explanation for explosive growth of the credit derivative tranche market.

We show that when the number of assets in the underlying portfolio of a tranche becomes large, the tranche converges to an option on the market portfolio. Specifically, if we restrict our attention to a tranche offering a digital payoff referenced to the loss on the underlying portfolio, the tranche payoff converges to the payoff of a digital market call option.[2] However, the previous section shows that holding the default probability constant, a digital call represents the cheapest to supply asset with a pre-specified credit rating. Because pooling and tranching synthetically creates the cheapest to supply asset in a given credit rating category, it effectively provides the optimal mechanism

[2] Although the digital payoff is a simplified version of tranche structures actively traded in real-world credit markets, this simplification is largely without loss of generality. To see this, it is sufficient to note that any non-digital tranche can be represented as a strictly positive combination of digital tranches. Consequently, the risk characterstics and pricing properties of a digital tranche carry over to the tranche structures traded in real-word credit markets.

for exploiting the arbitrage opportunity created by agents employing a naïve pricing model, which prices bonds solely on the basis of their expected payoff (i.e. credit rating). In other words, aside from completing the market by increasing the supply of highly-rated securities, the growth of the credit tranche market can potentially be explained as an endogenous, institutional response to an arbitrage opportunity in the credit markets.

To verify this claim we examine the risk characteristics and pricing of a prototypical tranche, offering a digital payoff referenced to the loss on the underlying portfolio of *economic assets*. Economic assets are assets whose conditional probability of default increases in the adversity of the economic state, and can be thought of in contrast to actuarial claims, whose default probability is unrelated to the economic state. In other words, economic assets are more likely to default in states of the world in which high marginal utility of wealth, or equivalently, in states where the value of one unit of consumption is high. For example, if we were to identify the priced states with realizations of the market return, economic assets could generally be described as assets whose expected value covaries positively with the realized market return (i.e. assets with a positive CAPM beta). This feature is typical of essentially all non-actuarial assets, and arises trivially in Merton's (1974) structural model of debt.

1.3 Integrating Merton's (1974) Credit Model with the CAPM

To fix ideas we examine the pricing and risk characteristics of a CDO tranche with an unconditional default probability, $\overline{p}_D$, written on a portfolio of economic assets, in this case - bonds. To build up an analytical model for pricing the CDO tranche we rely on the Merton structural model to determine the individual bond default probabilities, and then derive the distribution of portfolio losses using a central limit theorem argument. Specifically, we begin with the assumption that firm asset values are characterized by the following stochastic differential equation,

$$\frac{dA_i}{A_i} = (r_f + \beta \cdot \eta)dt + \beta\sigma_m dZ_m + \sigma_\epsilon dZ_i, \tag{10}$$

where r_f is the riskless rate, β is the CAPM beta of the asset returns on the market portfolio, η is the market risk premium and σ_m and σ_ϵ and the market and idiosyncratic asset return volatilities, respectively. We make the common assumption that a firm defaults whenever the terminal value of its assets, A_T, falls below the face value of debt, D. Crucially, we assume asset returns satisfy a CAPM relationship, which will later allow us to represent the CDO tranche as a derivative claim on the market portfolio. Using the distribution of asset returns conditional on the realization of the T-period market return, $r_{M,T}$, it is easy to show that the an individual firm's conditional probability of default is given by,

$$p_D(r_{M,T}) = \Phi\left[\frac{\ln\frac{D}{A_t} - \left(r_f + \beta\left(\frac{r_{M,T}}{T} - r_f\right) - \frac{\sigma_\epsilon^2}{2}\right)\cdot T}{\sigma_\epsilon\sqrt{T}}\right], \tag{11}$$

where the expression appearing in the brackets can be interpreted as the conditional distance to default given an observed market return of $r_{M,T}$. As posited earlier, the CAPM beta of economic assets is positive ($\beta > 0$), causing their conditional default probability to decrease with the magnitude of the T-period market return, $r_{M,T}$ ($\frac{dp_D(r_{M,T})}{dr_{M,T}} < 0$). Conveniently, after conditioning on the realization of the market return, asset returns and defaults are independent and idiosyncratic. This implies that the distribution of the number of defaulted firms in the underlying portfolio of bonds will be binomial with parameter $p_D(r_{M,T})$.

To derive the portfolio loss distribution we assume that the loss given default for each individual firm is a random variable between $[0, 1]$, with mean l and variance v^2. In other words, losses given default are purely idiosyncratic and have identical expectations under the objective and risk-neutral probability measures.[3] The conditional portfolio loss is then given by an equal-weighted sum of the issuer-specific losses:

$$\tilde{L}_p(r_{M,T}) = \frac{1}{N} \sum_{i=1}^{\tilde{N}(r_{M,T})} \tilde{L}_i(r_{M,T}) \tag{12}$$

Although it is typically not possible to derive a convenient analytical expression for the distribution of the conditional portfolio loss, $\tilde{L}_p(r_{M,T})$, progress can be made *via* a central limit theorem argument. When the number of firms in the underlying bond portfolio, N, is sufficiently large the central limit theorem indicates the portfolio loss distribution converges to,

$$\tilde{L}_p(r_{M,T}) \sim \phi\left(l \cdot p_D(r_{M,T}), \frac{p_D(r_{M,T}) \cdot \left((1 - p_D(r_{M,T})) \cdot l^2 + v^2\right)}{N}\right), \tag{13}$$

With the conditional portfolio loss distribution in hand we can now characterize the payoffs to various tranches written on the bond portfolio. In what follows, we focus on the pricing of a digital tranche, which pays one dollar when the (percentage) portfolio loss is less that X, and zero otherwise. The (percentage) magnitude of the portfolio loss beyond which the tranche ceases to pay, X, is known as the tranche attachment point.

Using the limiting distribution of the portfolio loss one can readily show that the conditional tranche default probability, is given by,

$$p_D(r_{M,T}, X, N) = 1 - \Phi\left(\frac{X - l \cdot p_D(r_{M,T})}{\sqrt{p_D(r_{M,T}) \cdot ((1 - p_D(r_{M,T})) \cdot l^2 + v^2)}} \cdot \sqrt{N}\right). \tag{14}$$

and the unconditional tranche default probability can be obtained by integrating the conditional default probability against the distribution of the T-period market return, $f(r_{M,T})$. The tranche attachment point is fixed such that the unconditional tranche default probability is equal to $\overline{p}_D$. Focusing on the asymptotic tranche behavior as the number of securities in the underlying portfolio

[3] The reduced-form credit literature typically fixes L^Q, and instead focuses on establishing the difference between the default intensity processes under the two measures, λ^P and λ^Q (see Duffie and Singleton (2003) and Pan and Singleton (2006)).

becomes large ($N \to \infty$),[4] the conditional tranche default probability converges to the following indicator function,

$$\lim_{N\to\infty} p_D(r_{M,T}, X, N) = \mathbf{1}_{X < l \cdot p_D(\hat{r}_{M,T})} \tag{15}$$

where $\hat{r}_{M,T}$ is implicitly defined through the condition, $X = l \cdot p_D(\hat{r}_{M,T})$. In fact, the binary nature of the conditional default probability indicates that the payoff function of the digital tranche converges (in probability) to the payoff function of a digital call option on the market portfolio. To see this more clearly, note that the tranche pays one dollar conditional on the market return being greater than $\hat{r}_{M,T}$ and zero otherwise. We formalize this result in the following theorem.

Theorem 1 *Suppose a digital tranche is written on a portfolio containing N identical economic assets, and has an attachment point of X, corresponding to an unconditional default probability of $\overline{p}_D$. As the number of securities in the portfolio underlying the tranche converges to infinity, $N \to \infty$, the tranche payoff function converges in probability to the payoff function of a digital market call with the same probability of expiring out of the money, and its price converges to the price of that call,*

$$\lim_{N\to\infty} P_0^{X,N} = e^{-r_f T} \cdot \Phi\left[\Phi^{-1}(1 - \overline{p}_D) - \frac{\eta}{\sigma_m} \cdot \sqrt{T}\right] \tag{16}$$

To obtain more intuition about the convergence of the tranche price to the price of the digital market call, $P_0^{X,\infty}$, as a function the number of securities in the underlying portfolio, N, we make use of the Arrow-Debreu pricing formalism. To do this, we specialize to the state-space defined by the realizations of the market return, $r_{M,T}$, and re-express the tranche price as an integral of the product of its state-contingent payoff expectation $(1 - p_D(r_{M,T}, X, N))$ with the state price $(q(r_{M,T}))$ across all possible states,

$$P_0^{X,N} = \int_{-\infty}^{\infty} (1 - p_D(r_{M,T}, X, N)) \cdot q(r_{M,T}) dr_{M,T} \tag{17}$$

Because the realization of $r_{M,T}$ orders states in ascending order of marginal utility and the derivative of the conditional tranche default probability with respect to the number of underlying securities, $\frac{\partial p_D(r_{M,T}, X, N)}{\partial N}$, is positive (negative) for market returns $r_{M,T}$ smaller (larger) than $\hat{r}_{M,T}$, increasing N effectively reallocates the probability of default from states with low marginal utility to states with high marginal utility. This causes the price of the digital tranche to decline monotonically in N, because it offers progressively less protection against systematically bad states. A full proof of this claim can be found in the Appendix.

Theorem 2 *The price of a digital tranche with attachment point, X, written on a portfolio of N identical assets, $P_0^{X,N}$, is monotonically decreasing in N, and converges to the price of the limiting*

[4] The focus on the asymptotic behavior is further warranted by the fact that our analysis already required N to be large to invoke the central limit theorem in the derivation of the portfolio loss distribution.

digital market call, $P_0^{X,\infty}$, as $N \to \infty$.

The strike price of the limiting digital market call option, can be obtained by setting the unconditional default probability of the digital tranche equal to p_D. Specifically, as $N \to \infty$, a simple manipulation shows that – in moneyness space – the option strike price is given by $K^* = \exp(\hat{r}_{M,T})$, where,

$$\hat{r}_{M,T} \quad = \quad \left(r_f + \eta - \frac{\sigma_m^2}{2}\right) \cdot T - \sigma_m \sqrt{T} \cdot \Phi^{-1}(1 - \overline{p}_D) \tag{18}$$

Finally, we assert without proof that the pricing results derived for the digital tranches carry over to the more standard tranches observed in practice, which are characterized by a pair of attachment points. To see this note that the payoff to a non-infinitesimally tight tranche, can be replicated by a *strictly positive* combination (i.e. portfolio) of digital tranches. Consequently, the pricing properties of the basis assets (i.e. the digital tranches) are inherited by the composite asset.

2 Data Description

Our empirical analysis relies on two main sets of data. The first consists of daily spreads of CDOs whose cash flows are tied to the DJ CDX North American Investment Grade Index. This index, which is described in detail in Longstaff and Rajan (2007), consists of a liquid basket of CDS contracts for 125 U.S. firms with investment grade corporate debt. Our data, which come from a proprietary database made available to us by Lehman Brothers, cover the period October 2004 to October 2006. The data contain daily spreads of the index as well as spreads on the 0-3, 3-7, 7-10, 10-15, and 15-30 tranches. As in Longstaff and Rajan (2007), we focus on the "on-the-run" indices which uses the first six months of CDX NA IG 4 through CDX NA IG 7 to produce a continuous series of spreads over the two-year period.

Our analysis also requires accurate prices for out-of-the-money market put options with five year maturity. During our sample period, no index options with maturity exceeding three years traded on centralized exchanges. However, two separate proprietary trading groups provided us with databases of daily over-the-counter quotes on five-year S&P 500 options. The two sources contain virtually identical quotes suggesting that the quotes reflect actual tradable spreads. The five-year option quotes include both at-the-money and 30 percent out-of-the-money put options which enable us to estimate a volatility skew for long-dated put options.

In addition to CDX and option data, we use a daily series of average corporate bond spreads on AA, A, BBB, BB, and B-rated bonds. These spreads, which were obtained from Lehman Brothers, are reported in terms of the 5-year CDS spread implied by corporate bond prices. Finally, we use the daily VIX obtained from the CBOE website. The VIX is a measure of near-term, at-the-money implied volatility of S&P 500 index options.

2.1 Summary Statistics

Table 1 provides summary statistics for the CDX index and tranche spreads as well as the bond spreads and the VIX series. Panel A reports average spread levels (and the VIX index level) and standard deviations for each of our series across the sample. As expected, average spreads are decreasing across the bond portfolios and across the tranches as the credit quality improves. The average spreads of the 3-7 and 7-10 tranches significantly exceed those of similarly rated bond portfolios across our sample period. However, both of these averages are strongly influenced by the early pricing of the CDX when, prior to it being widely accepted as a benchmark, mezzanine and senior tranche spreads were highly inflated. For example, as Figure 3 indicates, since October 2005 the 7-10 tranche spread has converged to that of the AA-rated bond portfolio. Indeed, tranche spreads have continued to match those of comparably rated bonds well into 2007.

Panels B reports daily correlations of each series in levels and Panel C reports correlations of first differences. In both levels and first differences, the AA are negatively correlated with the CDX index and its tranches. They are also negatively correlated in levels and in changes with the VIX volatility index. All other bonds are positively correlated with the CDX spread. The CDX and tranches have high correlations with each other and with the VIX, suggesting that market volatility is a key factor in the pricing of the CDX and its tranches but is not important in determining investment grade corporate bond yields.

3 Model Calibration

To calibrate the proposed structural model to empirical data on the credit default index (CDX) and its tranches, we rely on the assumption that the CDX is comprised of bonds issued by N identical firms. Historically, the representative firm included in the CDX index has had a credit rating of BBB or A.[5] We use this fact, combined with estimates of credit risk premia and the historical average CDX spread – approximately 45 bps over our sample period – to determine the implied bond recovery rate and the characteristics of the underlying firms (debt-to-asset ratios, asset betas, and asset idiosyncratic volatility) consistent with the Merton model estimates of the firms' default probabilities. We then use this information, combined with properties of the market return process (mean equity premium, market volatility), to derive the strike prices for the index puts spreads used to replicate the payoff to the CDO tranche.

3.1 Determining Firm Characteristics and Recovery Rates

The goal of first step of our calibration will be to match the ratio of the risk-neutral, λ^Q, and objective default intensities, λ^P, for the representative bond in the CDX to values reported in the credit literature, while simultaneously matching the mean historical yield spread on the CDX

[5] The decision regarding the credit rating of the representative firm is important, because it determines the "credit rating" of the index itself. Kakodar and Martin (2004) report that the CDX index had an average rating of BBB+ at the end of June 2004.

index. Because the ratio of λ^Q and λ^P reflects the relative importance of the risk-premium term (i.e. second term in (3)) in the pricing of a defaultable bond it is frequently considered a measure of credit risk. For example, if a bond's defaults are idiosyncratic and the recovery rate conditional on default is zero, the ratio is equal to one indicating that no additional risk premium is being attached to the timing of the defaults. Conversely, the higher a security's propensity to default only in states with high marginal utility the higher the value of the ratio. For example, Hull, Predescu and White (2005) report ratios of $\frac{\lambda^Q}{\lambda^P}$ of 9.8 for A-rated bonds and 5.1 for BBB-rated bonds. Consistent with intuition, this indicates that the average economic state in which an A-rated bond defaults is worse than the average state in which a BBB-rated security is likely to default. By matching this ratio we ensure that the component securities in the CDX have the correct quantity of systematic default risk relative to idiosyncratic default risk. On the other hand, by matching the historical yield spread on the CDX, we ensure that each component bond is priced correctly, and consequently, so is the index. To match the CDX yield spread, which itself is a function of the recovery rate, we will also need to develop the pricing formula for the credit index.

The risk-neutral default intensity for the index can be backed out from an estimate of the annualized CDX yield, y_{CDX}, and risk-neutral recovery rate, R^Q, through the following formula,

$$e^{-y_{CDX}\cdot T} = e^{-r_f\cdot T}\cdot\left(e^{-\lambda^Q\cdot T} + R^Q\cdot(1-e^{-\lambda^Q\cdot T})\right) \tag{19}$$

Formally, this equation states the the CDX price (left-hand side) is equal to the discounted value of the expected payoff under the risk-neutral measure (right-hand side).[6] Using a series of simple linear expansions it can be shown that the risk-neutral CDX default intensity, λ^Q, satisfying this condition is approximately equal to the ratio of the CDX yield spread and the (expected) loss given default, $\frac{y_{CDX}-r_f}{1-R}$.

The objective, or real-world, default intensity is simply the annualized rate of default on the bond. If the T-period default probability is $\overline{p}_D$, the corresponding default intensity is,

$$\lambda^P = -\frac{1}{T}\ln\left(1-\overline{p}_D\right)$$

In our modified version of the Merton model in which asset returns obey a CAPM relation, and a firm is considered in default only if its asset value at maturity, A_T, is below the face value of debt the unconditional default probability, $\overline{p}_D$, can be derived in closed form.[7] Specifically, we can show that the objective default intensity, λ^P, is determined by market volatility and a series of firm-specific parameters, which include the firm's debt-to-asset ratio, $\frac{D}{A_t}$, asset beta, β_a, and

[6] In what follows, we assume that the recovery rate is constant across all states, and consequently, the risk-neutral and objective measure recovery rates are identical. We therefore drop the superscript on the recovery rate. Note also that $(1-R)$ is the expected loss conditional on default denoted by l in Section 1.

[7] This assumption can be trivially modified to allow defaults at intermediate dates, for example if the asset value falls below the face value of debt at any point in time between t – the origination of the bond – and its maturity, T, as in Hull, Predescu, and White (2006).

idiosyncratic asset volatility, σ_ε, through the following expression,

$$\lambda^P = -\frac{1}{T}\ln\left(1-\Phi\left[\frac{\ln\frac{D}{A_t}-\left(r_f+\beta\eta-\frac{\beta_a^2\sigma_m^2+\sigma_\epsilon^2}{2}\right)\cdot T}{\sqrt{\beta_a^2\sigma_m^2+\sigma_\epsilon^2}\cdot\sqrt{T}}\right]\right) \tag{20}$$

where the expression inside $\Phi(\cdot)$ is the distance-to-default under the objective probability measure.

To price the CDX index we first derive the formula for its expected payoff, as a function of the market return. Because the *expected* payoff to the CDX index is identical to the *expected* payoff of an underlying bond, this step effectively corresponds to pricing the representative bond in the CDX. The expected payoff to the CDX is given by the sum of payoffs on the N bonds underlying the index,

$$\begin{aligned} E[CDX(r_{M,T})] &= E\left[\frac{1}{N}\sum_{i=1}^{N} 1\cdot\mathbf{1}_{A_{i,T}(r_{M,T})>D}+R\cdot\mathbf{1}_{A_{i,T}(r_{M,T})<D}\right] \\ &= 1-(1-R)\cdot p_D(r_{M,T}) \end{aligned} \tag{21}$$

To determine the price of the CDX index we simply apply the Arrow-Debreu valuation technique to the above conditional payoff expectation. By integrating the product of the conditional CDX payoff and the state price, $q(r_{M,T})$, across all realization of the market returns we obtain the price of the CDX (or, equivalently, the price of the representative bond),

$$P_{CDX}(R) \equiv \int_{-\infty}^{\infty} E[CDX(r_{M,T})]\cdot q(r_{M,T},\sigma_{IV})dr_{M,T}$$

In pricing the CDX index we exploit the fact that the state prices can be found in closed-form (see Breeden and Litzenberger (1978)) when the market follows a log-normal diffusion. Furthermore, by using the time-series average of the implied volatility for the 5-year index option as our estimate of σ_{IV}, this approach allows us to ensure that the pricing of the bonds underlying the CDX is roughly consistent with option prices. The spirit of this approach is similar to the recent work by Cremers et al. (2007), which shows that the pricing of individual credit default swaps is indeed consistent with the option-implied pricing kernel. Specifically, we set $\sigma_{IV} = 25\%$, and assume – for the purposes of the calibration – that the implied volatility does not change across option strikes. We then convert the CDX price into its equivalent yield spread (relative to the riskless rate) through, $-\frac{1}{T}\ln P_{CDX}(R) - r_f$, and match this quantity to the historical mean of the CDX spread (45 bps).

Consequently, the first part of the calibration procedure requires that we find a set of parameters $\{\frac{D}{A_t}, \beta_a, \sigma_\varepsilon, R\}$, which satisfies the restriction on the ratio of the risk-neutral and objective default intensities consistent with the average credit quality of the bonds underlying the index and matches the mean historical yield spread on the CDX index. In general, there may be multiple solutions to this system of non-linear equation given that we have only two constraints and four parameters. To ameliorate this problem we additionally require that the calibrated parameters roughly match the empirically observed characteristics of the underlying firms, such us equity volatility, equity beta and debt-to-asset ratio.

3.2 Determining the Strike Prices of the Replicating Put Spread

In the second step we solve for the strike prices of the index put spread, which replicates the tranche payoff. These strike prices are found by solving for the level of the market (i.e. moneyness) for which the *expected* CDX loss is equal to a given value, say $X\%$. The expected CDX loss is equal to $E[L_p(r_{M,T})] = 1 - E[CDX(r_{M,T})]$, or equivalently,

$$E[L_p(r_{M,T})] = (1-R) \cdot p_D(r_{M,T})$$

Setting this value equal to X and solving for $\exp(r_{M,T})$, yields the corresponding put price strike, K_X,[8]

$$K_X = \exp\left\{\frac{1}{\beta}\left(\ln\frac{D}{A_t} - \left(r_f \cdot (1-\beta) - \frac{\sigma_\epsilon^2}{2}\right)\cdot T - \sigma_\epsilon\sqrt{T}\cdot\Phi^{-1}\left(\frac{X}{1-R}\right)\right)\right\} \tag{22}$$

Repeating this procedure for the upper and lower tranche attachment points of the tranche yields the strike prices of the puts included in the replicating put spread. Consequently, the payoff to a tranche with a lower attachment point of X and upper attachment point of Y, is approximated by the payoff obtained by buying a riskless bond, writing a market index put at K_X, and buying a market index put with a strike price of K_Y.

Since the average bond in the CDX over our sample period is rated between BBB (BBB) and A, we calibrate the model under each of these two scenarios. The calibrated model parameters and the resulting index attachment points (put strike prices) are collected in Table 2.

3.3 Calculating Yields on Index Put Spreads

Having determined the relevant attachment points, or index put strike prices, that identify a portfolio that matches the systematic risk of various CDX tranches, we now want to be able to compare prices. To do so, we calculate the time series of yields for the index put spreads implied by the model. Each day the value of the replicating portfolio, V_t, is obtained by summing the value of a discount bond with face value of one and a maturity matching that of the CDX, a short position of q index put options struck at the lower loss attachment point (higher strike price, K_H), and a long position of q index put options at the upper loss attachment point (lower strike price, K_L). The quantity of options, $q = \frac{1}{K_H - K_L}$, is set such that the exposure to the market is eliminated outside of the range of strike prices (see the tranche payoff displayed in Figure 2), and remains constant over the life of the tranche. The yield of the replicating portfolio is simply $-\frac{1}{T}\ln V_t$.

The key input to this valuation procedure is the behavior of long-term *implied* volatility, which determines the pricing of the put options in the replicating portfolio. Specifically, because the puts

[8] All option strike prices are expressed as a fraction of the spot price.

present in the replicating portfolio will generally be deep out-of-the-money, both the at-the-money implied volatility and the implied volatility skew for 5-year index options will play a central role in determining the yield on the replicating portfolio. To ensure the robustness of our results we value the index put spread using two models of the implied volatility skew in long-term options. Using data on 5-year implied volatilities at-the-money and 30% out-of-the-money that are viewed to be tradeable by the financial intermediaries who supplied the data, we model the skew (1) by linearly extrapolating at the observed skew between a moneyness of 0.7 and 1; and (2) assuming no skew below a moneyness of 0.7 (i.e. fixing the implied volatility at moneyness levels smaller than 0.7 at the implied volatility of the 0.7 moneyness put option).

4 Empirical Results

Table 3 presents the calibration results. We compute the average spreads predicted by our model and compare these to the spreads offered by each of the CDX tranches. We also calculate the daily correlations between the two sets of spreads. Across all tranches, our model predicts significantly greater spreads than are present in the data. With a linear skew, the disparity gets worse (as a fraction of the spread) as the tranches increase in seniority. The 7-10 tranche spread predicted by our model exceeds actual spreads by more than a factor of two. For 10-15 and 15-30 tranches our model predicts spreads that are three times as large as in the data. Figures 4-7 graph the predicted and actual spreads through time. As can be seen, the model spreads significantly exceed actual spreads across the entire sample for each of the senior tranches. Only the 3-7 tranche is matched by our model at some point during the sample period. Moreover, because of the steady decline in senior tranche spreads over the sample period, by the end of the period the mispricing is even worse. As Figures 4-7 show, by the end of September 2006 the model spreads exceed actual spreads in the 7-10, 10-15, and 15-30 tranches by roughly a factor of five.

On the other hand, correlations in spread levels (daily) and changes between our model and observed spreads are uniformly large. This suggests that although their price levels are off by an order of magnitude, the returns offered by the put spread identified by our model and its corresponding CDX tranche are driven by common economic risk factors.

Table 3 also presents two robustness checks. The first examines the spread offered by the put spread if we assume no volatility skew on options with moneyness below 0.7. In this case, the option prices are computed using the implied volatility on the option that is 30% out-of-the-money. As we can see, this has a modest impact on the predicted spreads for all but the most senior tranche. In fact, for the 3-7 tranche, eliminating the skew increases the spread because the upper attachment point has a moneyness of 0.67 and so the main effect is to decrease the price of the lower put option (which is purchased by the investor). The 15-30 tranche has its predicted spread significantly reduced when the skew is eliminated. Because the upper attachment point has a moneyness of 0.37 its price is far higher when a linear skew is extrapolated from 0.7 to 0.37.

The second robustness check involves replacing the BBB bonds in our CDX with bonds rated single A. Interestingly, tranching a portfolio of bonds with higher average credit quality has the effect of increasing the senior tranche spreads. In other words, given the observed level of the CDX, improving the credit quality of the average bond increases the implied level of market risk to which each bond is exposed. As a result, the more senior tranches that default only on bad market outcomes must offer greater compensation to investors for bearing these risks.

Although the calibration offers an economically motivated and statistically sound pricing estimate, it does not offer much in the way of sensitivity analysis. An alternative approach is to ask what option strike prices are required to match observed CDX tranche spreads. Table 4 presents the annual yield spreads offered by put spreads with different upper and lower strike prices. Panel A presents yields calculated using a linearly extrapolated put spread. Panel B presents yields calculated with no skew in implied volatilities below moneyness of 0.7.

First, consider the 7-10 tranche, which offered a spread of 43.7 basis points over the sample period. Looking across the various put spreads, we see that 43 or more basis points are offered by put spreads with an upper strike price of 0.45 and a lower strike price of 0.4 or an upper strike price of 0.5 and a lower strike price of 0.3. Thus, to be willing to purchase the 7-10 tranche during our sample period, one must have believed that it was less likely to default than a 50 percent out-of-the-money 5-year put option on the S&P 500 was to expire in-the-money.

Given the decline in senior CDX tranche spreads across our sample, the yield spreads at the end of the sample paint an even worse picture. At the end of the sample, the 7-10 yield spread traded at 15 basis points. This figure can be compared to the yield spreads offered by put spreads at the end of the sample period. In Panel C, we see that the 40-30 put spread matches the yield spread of the 7-10 tranche at the end of the sample. An investor who purchases the 7-10 tranche must believe a decline of roughly 65% in the market over the next five years is more likely than a default of 15-20 percent of the investment-grade industrial firms in the US (assuming a 50 percent recovery rate).

5 Discussion

Our model provides a novel vantage point from which one can assess the recent developments in the structured finance domain, and suggests some directions for the future growth of this, already burgeoning, market segment. First, due to its focus on pricing, our model suggests a characterization of the equity tranche that is distinct from the conclusions offered by agency theory and asymmetric information (DeMarzo (2005)). Although, in the presence of asymmetric information about the cash flows of the underlying securities, the equity tranche is indeed very risky to the uninformed – as emphasized by DeMarzo (2005) – its cash flow risk is primarily of the idiosyncratic variety. In other words, because the equity tranche bears the first losses on the underlying portfolio, it is exposed primarily to diversifiable, idiosyncratic losses. The benign nature of the underlying risk – reflected in its low equilibrium price – stands in marked contrast to the tranche's popular characterization as "toxic waste." Although issuers of structured products are often required to hold this tranche as a means of alleviating the asymmetric information problem emphasized by DeMarzo (2005), they are also likely to be overcharging clients for this seemingly dangerous service.

Second, our theoretical derivations show that if the marginal investor prices a structured product solely on the basis of its credit rating (or, equivalently, default probability), the magnitude of the product's mispricing will grow with the number of securities, N, included in the underlying portfolio. Specifically, as the value of N becomes larger and the portfolio becomes more granular, tranches bear progressively more systematic risk, and should trade at higher yield spreads. If this is not the case, originators of structured products seeking to exploit this pricing error, will have a natural incentive to create products with more granular collateral pools ($N \to \infty$), e.g. comprised of loans, credit-card receivables, etc. The potential profitability of this scheme is further accentuated by a results of a recent investor survey conducted by the Bank of International Settlements (2005). The survey revealed that "there is much more appetite for granular than for non-granular pools," suggesting the presence of a natural clientele.

Finally, the deviation of tranche spreads from their default probabilities is predicted to be the greatest when the underlying securities have significant systematic risk themselves. This creates an incentive to supply structured products in which the underlying assets are themselves instruments with significant credit risk, i.e. high values of $\frac{\lambda^Q}{\lambda^P}$. As we have shown, senior tranches fit this description precisely, suggesting a potential explanation for the recent appearance of products such as the CDO^2, where the collateral pool is comprised of various CDO tranches, and $CPDOs$, which provide leveraged exposures to highly-rated credit portfolios. Because the very senior tranches have tiny unconditional probabilities of default, while the highest credit rating is AAA, the suppliers or originators of CDOs are leaving money on the table unless they lever up these securities to match more closely the default probabilities of other AAA-rated securities.

6 Conclusion

This paper presents a framework for understanding the risk and pricing implications of structured finance activities. We demonstrate that senior CDO tranches have significantly different systematic risk exposures than their credit rating matched counterpart, and should therefore command different risk premia. Credit rating agencies, including sophisticated services like KMV, do not provide customers with adequate information for pricing. Forecasts of unconditional cash flows are insufficient for determining the discount rate and therefore can create significant mispricing in derivatives on bond portfolios.

In the spirit of Arrow-Debreu, we develop an intuitive approach relative to the current statistics-heavy approaches that are popular in the credit literature and in practice. Our pricing strategy for credit default obligations (CDOs) and the claims, or tranches, written against them is to identifying packages of other investable securities that deliver identical payoffs conditional on the market. Projecting expected cash flows into market return space may be an effective way to identify investable portfolios that replicate the systematic risk in other applications. Our analysis demonstrates that an Arrow-Debreu approach to pricing can be operationalized relatively easily.

Our pricing estimates suggest that investors in senior CDO tranches are grossly undercompensated for the highly systematic nature of the risks they bear. We demonstrate that an investor willing to incur these risks bears highly similar risks by writing out-of-the-money put options on the market but receives far greater compensation for doing so. We argue that this discrepancy has much to do with the fact that credit rating agencies are willing to certify CDO tranches as “safe” when, from an asset pricing perspective, they are quite the opposite.

References

[1] Black, Fischer and Myron Scholes, 1973, The pricing of options and corporate liabilities, *Journal of Political Economy* 81, 617-654.

[2] Breeden, Douglas, and Robert Litzenberger, 1978, Prices of state-contingent claims implicit in option prices, *Journal of Business*, 621-651.

[3] Cremers, Martijn, Joost Driessen, and Pascal Maenhout, 2007, Explaining the Level of Credit Spreads: Option-Implied Jump Risk Premia in a Firm Value Model, *Review of Financial Studies* (forthcoming).

[4] DeMarzo, Peter, 2005, The pooling and tranching of securities: A model of informed intermediation, *Review of Financial Studies*, 1-35.

[5] Duffie, Darrell, and Kenneth Singleton, 2003, *Credit Risk*, Princeton University Press.

[6] Hull, John, Mirela Predescu, and Alan White, 2005, Bond prices, default probabilities, and risk premiums, *Journal of Credit Risk*, 53-60.

[7] Hull, John, Mirela Predescu, and Alan White, 2006, The Valuation of Correlation-dependent Credit Derivatives Using a Structural Model, *mimeo*.

[8] Kakodkar, Atish and Barnaby Martin, 2004, The Standardized Tranche Market Has Come Out of Hiding, *Journal of Structured Finance*, Fall, 76-81.

[9] Lintner, J., 1965, The valuation of risk assets and the selection of risky investments in stock portfolios and capital budgets, *Review of Economics and Statistics*, 13-37.

[10] Longstaff, Francis, and Arvind Rajan, 2007, An Empirical Analysis of the Pricing of Collateralized Debt Obligations, UCLA working paper.

[11] Merton, Robert C., 1973, Theory of rational option pricing, *Bell Journal of Economics and Management Science*, 141-183.

[12] Merton, Robert C., 1974, On the pricing of corporate debt: The risk structure of interest rates, *Journal of Finance*, 449-470.

[13] Pan, Jun, and Kenneth Singleton, 2006, Default and recovery implicit in the term structure of sovereign CDS spreads, MIT working paper.

[14] Sharpe, William F., 1963, A simplified model for portfolio analysis, *Management Science*, 277-293.

A Properties of the Digital Tranche

If we let $f(r_{M,T})$ and $F(r_{M,T})$ denote the probability density function and cumulative density function of the T-period market return, the unconditional default probability of the digital tranche is given by,

$$\begin{aligned} p_D(X,N) &= \int_{-\infty}^{\infty} p_D(r_{M,T},X,N)\cdot f(r_{M,T})dr_{M,T} \\ &\to \int_{-\infty}^{\hat{r}_{M,T}} f(r_{M,T})dr_{M,T} = F(\hat{r}_{M,T}) \quad \text{as} \quad N\to\infty \end{aligned} \tag{23}$$

Setting this expression equal to the desired unconditional default probability, $\overline{p}_D$, then allows us to obtain the strike price of the digital market call option to which the tranche converges as $N \to \infty$. The option strike price is given by $K^* = \exp(\hat{r}_{M,T})$, where,

$$\hat{r}_{M,T} = \left(r_f + \eta - \frac{\sigma_m^2}{2}\right)\cdot T - \sigma_m\sqrt{T}\cdot\Phi^{-1}(1-\overline{p}_D) \tag{24}$$

To price the digital tranche we define a state-space over the realizations of the market return, $r_{M,T}$, and make use of the Arrow-Debreu pricing formalism. In other words, we price the tranche by summing the product of its state-contingent payoff expectation with the state price across all possible states. If we let $p_D(r_{M,T},X,N)$ denote the state-contingent default probability for the digital tranche, and $q(r_{M,T})$, denote the Arrow-Debreu state price for state, $r_{M,T}$, the tranche price will be given by,

$$P_0^{X,N} = \int_{-\infty}^{\infty}(1-p_D(r_{M,T},X,N))\cdot q(r_{M,T})dr_{M,T} \tag{25}$$

Substituting the asymptotic tranche default probability into (17) yields the following expression for the tranche price,

$$\begin{aligned} P_0^{X,\infty} &= \int_{-\infty}^{\infty}(1-\mathbf{1}_{X\,<\,l\cdot p_D(\hat{r}_{M,T})})\cdot q(r_{M,T})dr_{M,T} \\ &= \int_{\hat{r}_{M,T}}^{\infty} q(r_{M,T})dr_{M,T} \end{aligned} \tag{26}$$

which, unsurprisingly, is the equation defining the price of a T-period digital call option on the market portfolio, with a strike price of $K^* = \exp(\hat{r}_{M,T})$ (in moneyness space).

Finally, the law of one price ensures that the value of the digital tranche is simply given by the value of the digital market call with a strike price of $K^* = \exp(\hat{r}_{M,T})$ (in moneyness space). Substituting in the expression for K^* into the Black-Scholes formula for the price of the digital market call we immediately obtain the price of the digital tranche,

$$P_0^{X,\infty} = e^{-r_f T}\cdot\Phi\left[\Phi^{-1}(1-p_D) - \frac{\eta}{\sigma_m}\cdot\sqrt{T}\right] \tag{27}$$

The tranche price can also be expressed in terms of a yield spread, λ^Q in excess of the riskless rate by setting $P_0^{X,\infty} = e^{-(r_f+\lambda^Q)\cdot T}$,

$$\lambda^Q = -\frac{1}{T}\ln\Phi\left(\Phi^{-1}(1-p_D)) - \frac{\eta}{\sigma_m}\cdot\sqrt{T}\right) \tag{28}$$

λ^Q has the interpretation of a risk-neutral default intensity and depends on the the expected payoff, as determined by the probability of default, and - on a risk adjustment determined by the T-period market Sharpe ratio. As intuition would suggest, whenever the market risk premium, η, is equal to zero, the tranche will trade at a yield spread equal to its objective default intensity λ^P.

A.1 Tranche Price Convergence as $N \to \infty$

The price of the digital tranche satisfies the following integral equation, expressing its price as the integral of its state contingent payoff and the corresponding state price,

$$P_0^{X,N} = \int_{-\infty}^{\infty} (1 - p_D(r_{M,T}, X, N)) \cdot q(r_{M,T}) dr_{M,T} \tag{29}$$

We are interested in establishing the direction and monotonicity of the convergence of the tranche price as $N \to \infty$. To do this we begin by noting that all tranches along the path $N \to \infty$ are constrained to have the same unconditional default probability $\overline{p}_D$. This implies that,

$$\frac{\partial}{\partial N}\int_{-\infty}^{\infty} p_D(r_{M,T}, X, N) f(r_{M,T}) dr_{M,T} = 0 \tag{30}$$

where $f(r_{M,T})$ is the probability density function of the T-period market return. Because increasing N has no effect on the unconditional tranche default probability, it also has no effect on the unconditional expectation of the tranche payoff, which is equal to $1 - \overline{p}_D$. In other words, two digital tranches with different values of N but identical unconditional default probabilities can be thought of as mean-preserving spreads of each other.

Now consider differentiating the expression for the tranche price with respect to N,

$$\frac{\partial P_0^{X,N}}{\partial N} = -\int_{-\infty}^{\infty} \frac{\partial p_D(r_{M,T}, X, N)}{\partial N} \cdot q(r_{M,T}) dr_{M,T}$$

where the partial derivative of the conditional tranche default probability changes sign at $\hat{r}_{M,T}$:

$$\frac{\partial p_D(r_{M,T}, X, N)}{\partial N} > 0 \quad \text{for} \quad r_{M,T} < \hat{r}_{M,T}$$
$$\frac{\partial p_D(r_{M,T}, X, N)}{\partial N} < 0 \quad \text{for} \quad r_{M,T} > \hat{r}_{M,T}$$

Multiplying and dividing each term in the expression for $\frac{\partial P_0^{X,N}}{\partial N}$ by the probability density function of the T-period market return, $f(r_{M,T})$, we obtain,

$$\frac{\partial P_0^{X,N}}{\partial N} = -\int_{-\infty}^{\infty} \frac{\partial p_D(r_{M,T}, X, N)}{\partial N} \cdot \frac{q(r_{M,T})}{f(r_{M,T})} \cdot f(r_{M,T}) dr_{M,T} \tag{31}$$

where the ratio, $m(r_{M,T}) = \frac{q(r_{M,T})}{f(r_{M,T})}$, represents the stochastic discount factor for state $r_{M,T}$, and is monotonically declining in $r_{M,T}$. To exploit the monotonicity property of the stochastic discount factor we first split the above integral at $\hat{r}_{M,T}$ and then apply the mean-value theorem twice,

$$\begin{aligned}\frac{\partial P_0^{X,N}}{\partial N} &= -\Big(\int_{-\infty}^{\hat{r}_{M,T}} \frac{\partial p_D(r_{M,T},X,N)}{\partial N} \cdot m(r_{M,T}) \cdot f(r_{M,T}) dr_{M,T} + \\ &\quad + \int_{\hat{r}_{M,T}}^{\infty} \frac{\partial p_D(r_{M,T},X,N)}{\partial N} \cdot m(r_{M,T}) \cdot f(r_{M,T}) dr_{M,T}\Big) \\ &= -\left(\overline{m} \cdot \int_{-\infty}^{\hat{r}_{M,T}} \frac{\partial p_D(r_{M,T},X,N)}{\partial N} \cdot f(r_{M,T}) dr_{M,T} + \underline{m} \cdot \int_{\hat{r}_{M,T}}^{\infty} \frac{\partial p_D(r_{M,T},X,N)}{\partial N} \cdot f(r_{M,T}) dr_{M,T}\right)\end{aligned}$$

where $\overline{m}$ is some number in $[m(-\infty), m(\hat{r}_{M,T})]$, $\underline{m}$ is some number in $[m(\hat{r}_{M,T}), m(\infty)]$. However, because $m(r_{M,T})$ is monotonic in the market return we immediately have $\overline{m} > \underline{m}$. Moreover, we have also established that for all N, the expectation of the derivative of the conditional default probability with respect to N evaluated under the objective probability measure satisfies,

$$\int_{-\infty}^{\hat{r}_{M,T}} \frac{\partial p_D(r_{M,T},X,N)}{\partial N} \cdot f(r_{M,T}) dr_{M,T} + \int_{\hat{r}_{M,T}}^{\infty} \frac{\partial p_D(r_{M,T},X,N)}{\partial N} \cdot f(r_{M,T}) dr_{M,T} = 0$$

In other words, the two terms in the above expression are of equal magnitude, ξ, but opposing sign. Consequently, we can conclude that,

$$\frac{\partial P_0^{X,N}}{\partial N} = -(\overline{m} \cdot \xi - \underline{m} \cdot \xi) = \xi \cdot (\underline{m} - \overline{m}) < 0$$

which indicates that the price of the digital tranche is monotonically declining in N, since the sign of the derivative does not depend on N.

Table 1
Summary Statistics for US Credit Market (9/2004 – 9/2006)

This table reports summary statistics of various credit market securities. The credit spreads for various bond indices correspond to the difference between 5-year S&P US bond yields and the 5-year swap rate. The CDX series is the Lehman Brothers investment grade index of 5-year US industrials credit default swaps. The CDX tranche spreads are denoted as [lower loss attachment, upper loss attachment]. The VIX is the Chicago Board Options Exchange (CBOE) volatility index measuring near term implied volatility from S&P500 index options. Five-year at-the-money implied volatility is denoted as 5-yr Vol.

Panel A: Time series means and standard deviations of daily series in basis points (volatility in percent)

	VIX	5yVOL	AA	A	BBB	BB	B	CDX	3-7	7-10	10-15	15-30
Mean	13.2	17.8	15.3	25.6	46.9	143.3	254.8	47.8	146.4	44.1	19.9	9.2
Std	1.8	1.3	2.2	2.6	5.6	22.8	36.7	6.9	56.8	24.9	9.0	3.7

Panel B: Correlations between daily series

	VIX	5yVOL	AA	A	BBB	BB	B	CDX	3-7	7-10	10-15	15-30
VIX	1.00											
5yVOL	0.36	1.00										
AA	0.01	0.55	1.00									
A	-0.03	0.13	0.24	1.00								
BBB	0.11	0.52	0.50	0.85	1.00							
BB	0.24	0.62	0.26	0.60	0.74	1.00						
B	-0.02	0.39	0.21	0.74	0.77	0.86	1.00					
CDX	0.08	-0.12	-0.24	0.57	0.43	0.56	0.68	1.00				
3-7	0.21	-0.41	-0.48	0.35	0.16	0.09	0.22	0.66	1.00			
7-10	0.16	-0.48	-0.43	0.30	0.13	-0.07	0.07	0.53	0.95	1.00		
10-15	0.14	-0.44	-0.43	0.42	0.23	0.11	0.26	0.72	0.96	0.95	1.00	
15-30	-0.02	-0.31	-0.41	0.54	0.30	0.36	0.50	0.87	0.70	0.61	0.79	1.00

Panel C: Correlations between first-differences of daily series

	VIX	5yVOL	AA	A	BBB	BB	B	CDX	3-7	7-10	10-15	15-30
VIX	1.00											
5yVOL	0.18	1.00										
AA	0.02	0.02	1.00									
A	0.10	0.06	0.66	1.00								
BBB	0.07	0.14	0.59	0.77	1.00							
BB	0.13	0.10	0.19	0.31	0.36	1.00						
B	0.12	0.09	0.08	0.21	0.30	0.63	1.00					
CDX	0.39	0.20	0.00	0.07	0.12	0.29	0.19	1.00				
3-7	0.14	0.08	0.04	0.13	0.21	0.26	0.23	0.20	1.00			
7-10	0.16	0.05	0.04	0.11	0.18	0.23	0.20	0.18	0.89	1.00		
10-15	0.17	0.09	0.07	0.12	0.20	0.27	0.23	0.25	0.84	0.86	1.00	
15-30	0.13	0.07	0.07	0.10	0.15	0.19	0.14	0.14	0.73	0.72	0.78	1.00

Table 2
Model Parameters and Resulting Attachment Points

	A	BBB	Source
RN / RW Intensity	9.8	5.1	Hull, Predescu, and White (2005)
Rf	4.5%	4.5%	Assumption
Market Risk Premium	5.0%	5.0%	Assumption
Market Vol	15.0%	15.0%	Assumption
Kernel Vol	25.0%	25.0%	Assumption
T	5.0	5.0	Contract maturity
Debt-to-Assets	0.55	0.55	Calibration
Asset Beta	0.54	0.54	Calibration
Idiosyncratic Vol	11.9%	13.1%	Calibration
Recovery Rate	42.3%	54.1%	Calibration
Attachment Points			
3.0%	0.6303	0.6733	Model output
7.0%	0.5164	0.5132	Model output
10.0%	0.4665	0.4478	Model output
15.0%	0.4068	0.3740	Model output
30.0%	0.2843	0.2421	Model output

Table 3
Comparison of Actual and Model Implied Tranche Spreads

The actual CDX tranche spreads corresponds to various tranches on the Lehman Brothers investment grade index of 5-year US industrials credit default swaps. The yield on the S&P 500 put spread is calculated using strike prices identified from a Merton credit model calibration procedure. The calibration assumes the CDX consists of 125 BBB bonds and a linear skew in 5-year implied volatilities.

Calibration assumes CDX consists of 125 BBB bonds

Tranche	Mean Model Spread [bps]	Mean Actual Spread [bps]	Correlation of Model and Actual	Correlation of Model and Actual (Changes)
Linear Skew				
3%-7%	187	146	0.68	0.18
7%-10%	111	44	0.51	0.18
10%-15%	74	20	0.54	0.23
15%-30%	34	9	0.43	0.17
No Skew below 0.7				
3%-7%	221	146	0.65	0.18
7%-10%	90	44	0.47	0.17
10%-15%	43	20	0.49	0.22
15%-30%	10	9	0.33	0.16

Calibration assumes CDX consists of 125 A bonds

Tranche	Mean Model Spread [bps]	Mean Actual Spread [bps]	Correlation of Model and Actual	Correlation of Model and Actual (Changes)
Linear Skew				
3%-7%	172	146	0.66	0.18
7%-10%	118	44	0.52	0.18
10%-15%	87	20	0.56	0.23
15%-30%	47	9	0.45	0.17
No Skew below 0.7				
3%-7%	192	146	0.63	0.18
7%-10%	99	44	0.48	0.17
10%-15%	57	20	0.51	0.23
15%-30%	18	9	0.36	0.16

Table 4
Yields on Various S&P 500 Put Spreads

The yield on the 5-year S&P 500 put spread is calculated using for a variety of strike prices. The put spread consists of shorting an index put option at the upper strike price and buying an index put option at the lower strike price. Panel A assumes a linear skew in long-term implied volatilities consistent with the observed skew between strike prices of 70 and 100. Panel B assumes that there is no skew below strike prices of 70.

Average Daily Yield from 9/2004 to 9/2006

Panel A: Linearly extrapolated skew in implied volatilities

Lower Strike Price	Upper Strike Price: 0.30	0.35	0.40	0.45	0.50	0.55	0.60	0.65	0.70	0.75
0.00	7	11	17	24	32	42	52	64	77	91
0.05	8	13	19	27	36	46	57	69	83	98
0.10	10	16	23	31	40	51	63	76	90	105
0.15	13	19	27	36	46	57	70	84	98	114
0.20	17	24	33	42	53	65	78	92	108	125
0.25	23	31	40	50	61	74	88	103	119	136
0.30		38	48	59	71	84	99	114	131	149
0.35			58	69	82	96	111	127	145	163
0.40				81	94	109	125	142	160	179
0.45					108	123	140	157	176	196
0.50						139	156	174	193	214
0.55							173	192	212	233
0.60								211	232	254
0.65									253	275
0.70										298

Panel B: No skew in implied volatilities below the strike price of 70

Lower Strike Price	Upper Strike Price: 0.30	0.35	0.40	0.45	0.50	0.55	0.60	0.65	0.70	0.75
0.00	1	2	5	10	18	28	41	57	77	100
0.05	1	3	6	12	20	31	45	62	83	107
0.10	1	3	7	13	22	34	49	68	90	115
0.15	2	4	9	16	25	38	55	75	98	125
0.20	3	6	11	19	30	44	62	83	109	137
0.25	4	8	14	23	35	51	71	94	121	151
0.30		12	19	29	43	61	82	107	136	168
0.35			26	38	54	73	96	123	154	189
0.40				50	67	89	114	143	176	213
0.45					85	108	136	167	203	242
0.50						132	162	196	233	274
0.55							192	228	268	312
0.60								265	307	353
0.65									351	399
0.70										448

Table 4 (Continued)

Yield on 9/11/2006

Panel C: Linearly extrapolated skew in implied volatilities

Lower Strike Price	Upper Strike Price: 0.30	0.35	0.40	0.45	0.50	0.55	0.60	0.65	0.70	0.75
0.00	5	9	14	20	27	35	44	55	66	79
0.05	6	10	16	22	30	39	49	60	72	85
0.10	8	12	18	25	34	43	53	65	78	91
0.15	10	15	22	29	38	48	59	72	85	99
0.20	13	19	26	35	44	55	66	79	93	108
0.25	18	24	32	41	51	63	75	88	103	118
0.30		31	39	49	60	72	84	98	113	129
0.35			48	58	69	82	95	110	125	142
0.40				68	80	93	107	123	139	156
0.45					92	106	121	136	153	171
0.50						120	135	151	168	187
0.55							150	167	185	204
0.60								184	203	222
0.65									221	241
0.70										262

Panel B: No skew in implied volatilities below the strike price of 70

Lower Strike Price	Upper Strike Price: 0.30	0.35	0.40	0.45	0.50	0.55	0.60	0.65	0.70	0.75
0.00	1	2	4	9	15	23	35	49	66	86
0.05	1	2	5	10	16	26	38	53	72	93
0.10	1	3	6	11	19	29	42	58	78	100
0.15	1	3	7	13	21	32	47	64	85	109
0.20	2	5	9	15	25	37	53	71	93	119
0.25	3	6	12	19	30	43	60	80	104	131
0.30		10	16	24	36	51	70	92	117	145
0.35			22	32	45	62	82	105	133	163
0.40				42	57	75	97	123	152	184
0.45					72	92	116	143	174	209
0.50						112	138	168	201	237
0.55							164	196	231	269
0.60								228	265	305
0.65									303	345
0.70										388

Credit Market Summary on 9/11/2006 in basis points (volatility in percent)

	VIX	5yVOL	AA	A	BBB	BB	B	CDX	3-7	7-10	10-15	15-30
Mean	13.0	19.3	15.9	28.1	52.5	158.9	264.1	37.3	71.5	15.5	7.5	4.0

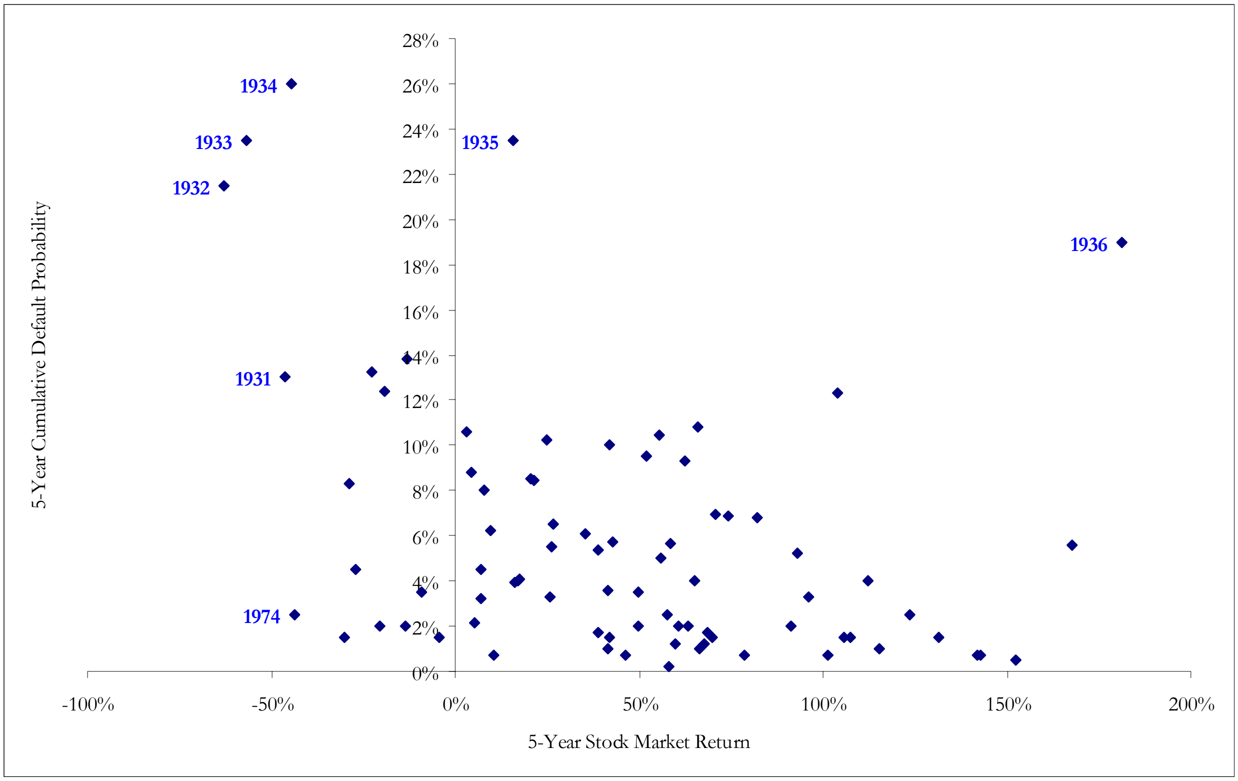

Figure 1. Historical 5-year corporate default rates conditional on the 5-year stock market return (1926-2005). Each year, the 5 most recent annual default rates from Moody's are summed and compared to the recent 5-year stock market return. The stock market return is the return on the CRSP value-weight market portfolio.

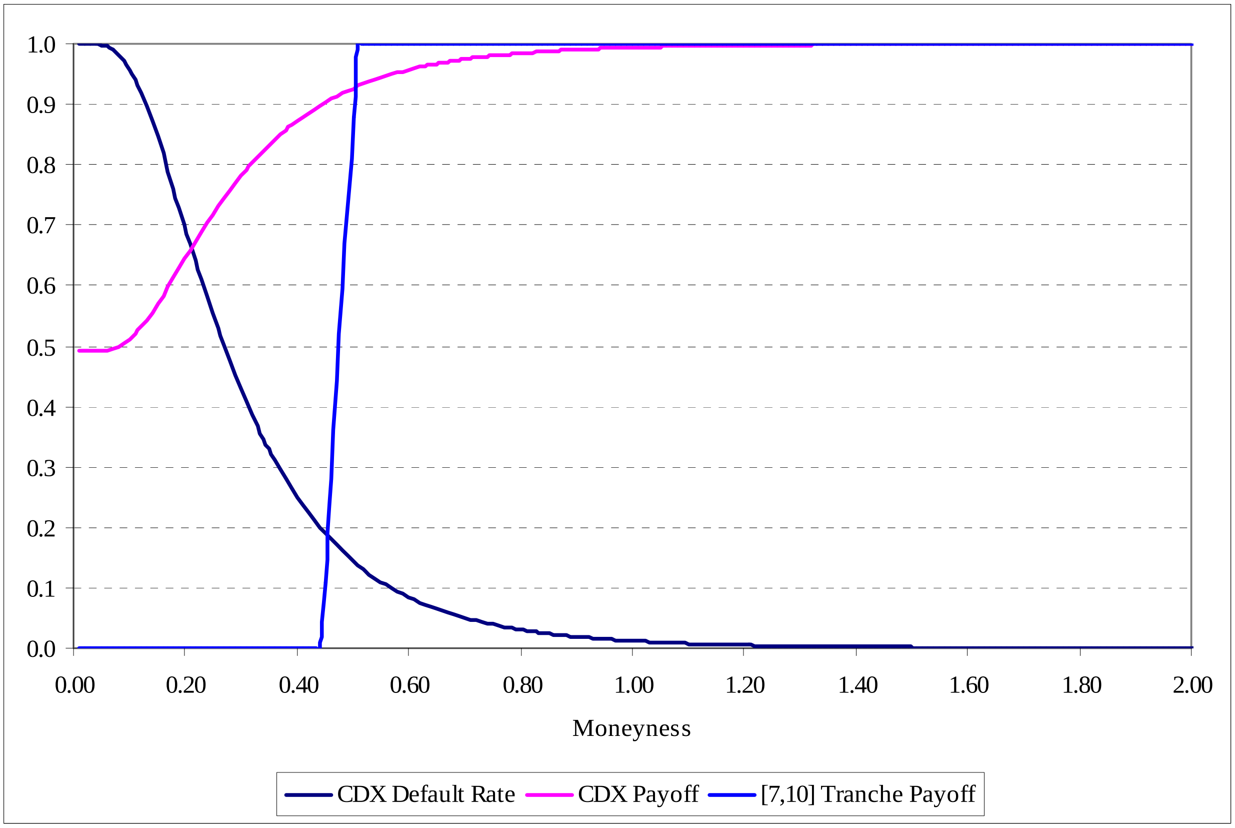

Figure 2. State Contingent CDX default rate, CDX payoff, and 7%-10% CDX tranche payoff.
The Merton model applied to CAPM assets is used to determine the CDX default rate conditional on the market return. The CDX is assumed to consist of 125 identical BBB bonds with a debt-to-asset ratio of 0.46, asset beta of 0.61, unconditional default intensity of 9bps, idiosyncratic volatility of 15.7%, and a recovery rate of 48%. The market portfolio has an annual volatility of 15%.

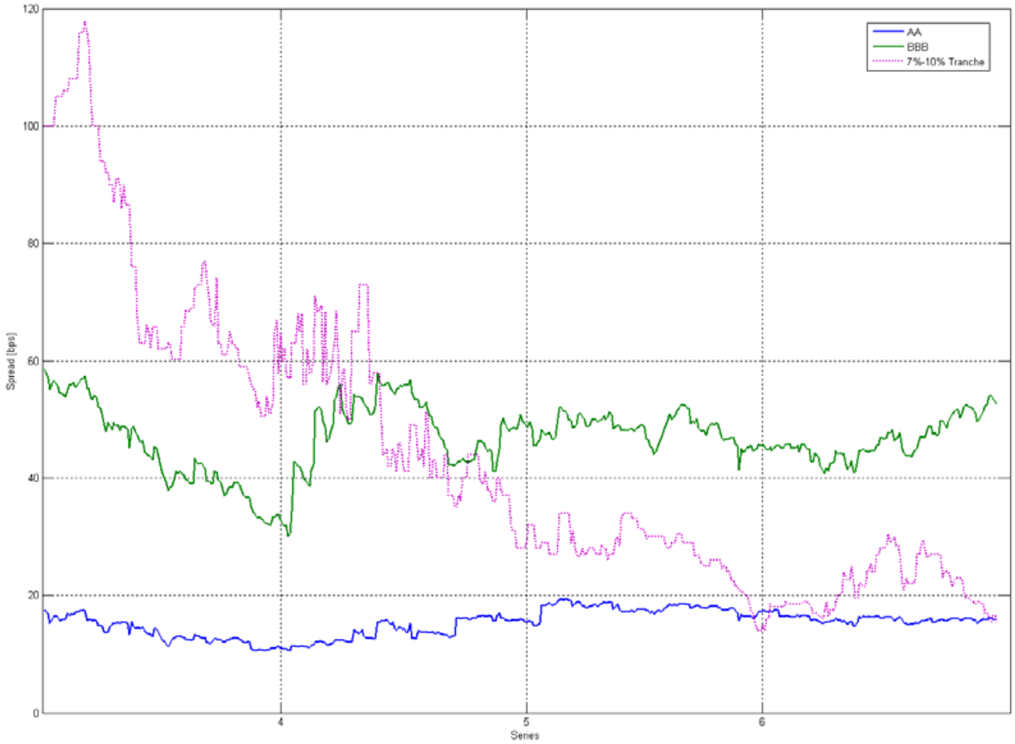

Figure 3. Time series of tranche spread and credit spreads for bond indices (9/2004 – 9/2006).
The credit spreads for the AA and BBB bond indices correspond to the difference between 5-year S&P US bond yields and the 5-year swap rate. The CDX tranche spread corresponds to the 7%-10% tranche on the Lehman Brothers investment grade index of 5-year US industrials credit default swaps.

Figure 4. Time series of tranche spread and S&P 500 index put spread (9/2004 – 9/2006).
The CDX tranche spread corresponds to the 3%-7% tranche on the Lehman Brothers investment grade index of 5-year US industrials credit default swaps. The yield on the S&P 500 put spread is calculated using strike prices identified from the calibration procedure. The calibration assumes the CDX consists of 125 BBB bonds and a linear skew in 5-year implied volatilities.

Figure 5. Time series of tranche spread and S&P 500 index put spread (9/2004 – 9/2006).
The CDX tranche spread corresponds to the 7%-10% tranche on the Lehman Brothers investment grade index of 5-year US industrials credit default swaps. The yield on the S&P 500 put spread is calculated using strike prices identified from the calibration procedure. The calibration assumes the CDX consists of 125 BBB bonds and a linear skew in 5-year implied volatilities.

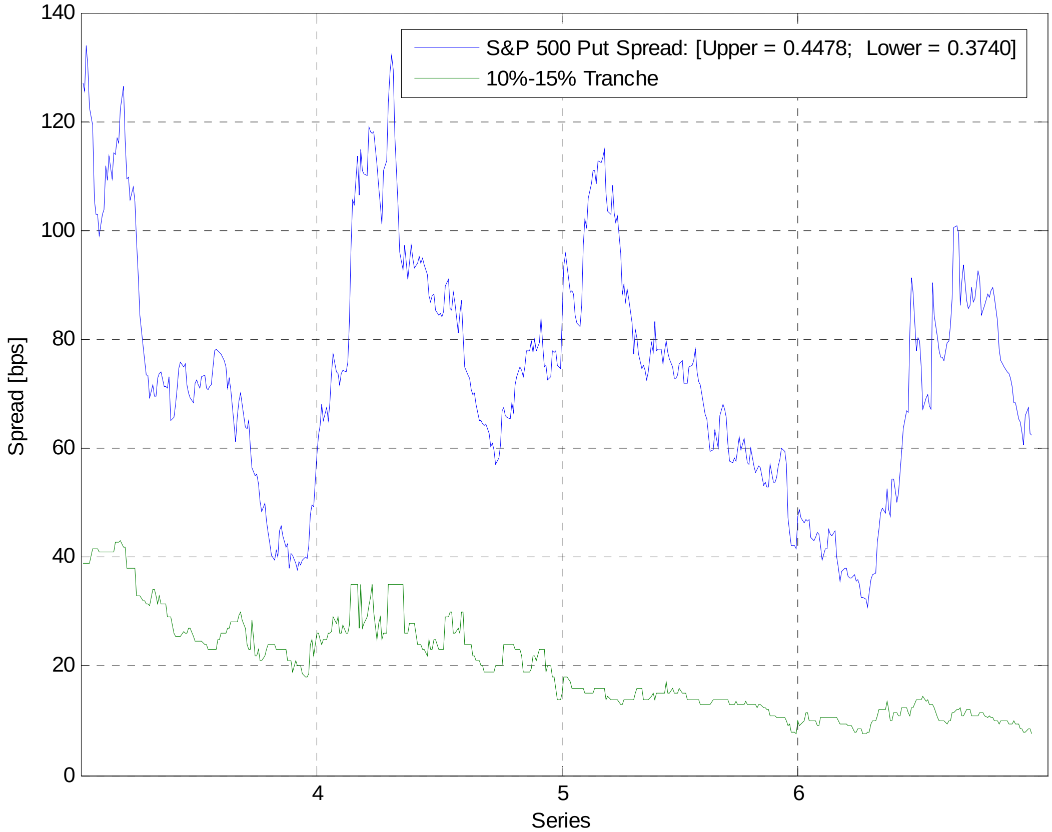

Figure 6. Time series of tranche spread and S&P 500 index put spread (9/2004 – 9/2006).
The CDX tranche spread corresponds to the 10%-15% tranche on the Lehman Brothers investment grade index of 5-year US industrials credit default swaps. The yield on the S&P 500 put spread is calculated using strike prices identified from the calibration procedure. The calibration assumes the CDX consists of 125 BBB bonds and a linear skew in 5-year implied volatilities.

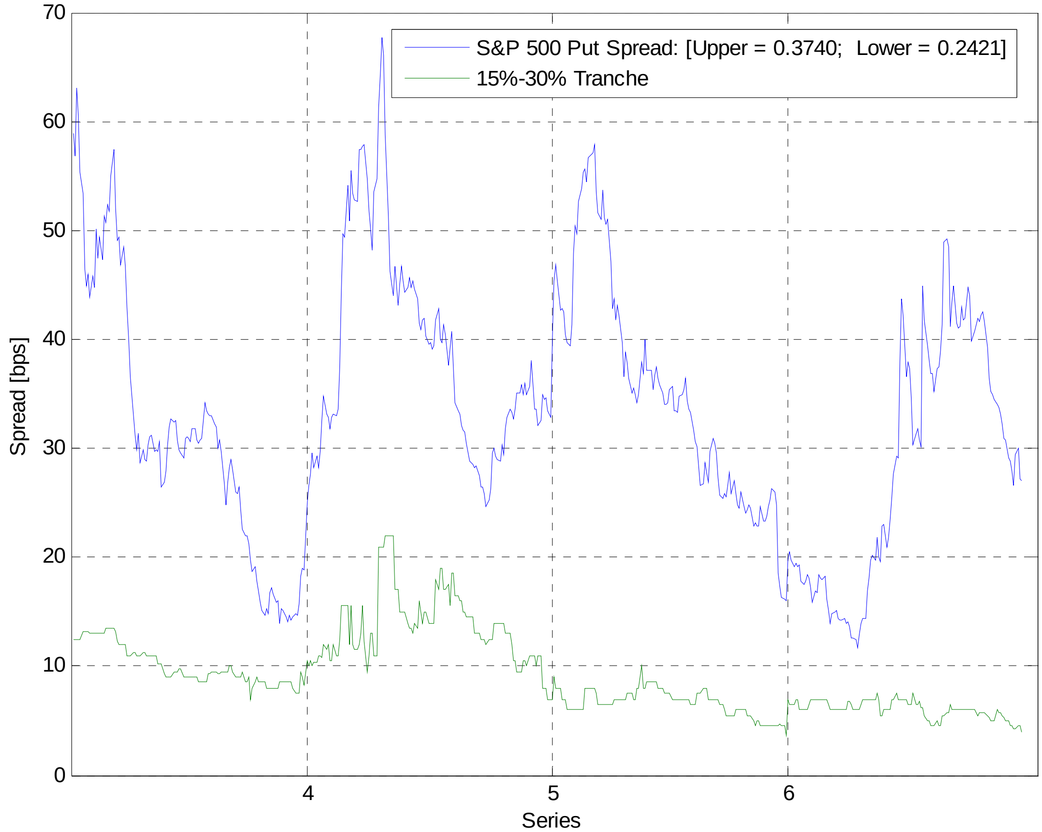

Figure 7. Time series of tranche spread and S&P 500 index put spread (9/2004 – 9/2006).
The CDX tranche spread corresponds to the 15%-30% tranche on the Lehman Brothers investment grade index of 5-year US industrials credit default swaps. The yield on the S&P 500 put spread is calculated using strike prices identified from the calibration procedure. The calibration assumes the CDX consists of 125 BBB bonds and a linear skew in 5-year implied volatilities.