

Economic Catastrophe Bonds

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

June 2007

Abstract

The central insight of asset pricing is that a security's value depends on both its distribution of payoffs across economic states and state prices. In fixed income markets, many investors focus exclusively on estimates of expected payoffs, such as credit ratings, 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. We show that many structured finance instruments can be characterized as economic catastrophe bonds, but offer far less compensation than alternatives with comparable payoff profiles. We argue that this difference arises from the willingness of rating agencies to certify structured products with a low default likelihood as “safe” and from a 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, Boston University, NYU, Virginia Tech, 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.

This paper investigates the pricing and risks of instruments created as a result of recent structured finance activities. Pooling economic assets into large portfolios and tranching them into sequential cash flow claims has become a big business, generating record profits for both the Wall Street originators and the agencies that rate these securities. A typical tranching scheme involves prioritizing the cash flows (liabilities) of the underlying collateral pool, such that a senior claim suffers losses only after the principal of the subordinate tranches has been exhausted. This prioritization rule allows senior tranches to have low default probabilities, garnering high credit ratings. However, it also confines senior tranche losses to systematically bad economic states, effectively creating economic catastrophe bonds.

The fundamental asset pricing insight of Arrow (1964) and Debreu (1959) is that an asset's value is determined by both its distribution of payoffs across economic states and state prices. Securities that fail to deliver their promised payments in the "worst" economic states will have low values, because these are precisely the states where a dollar is most valuable. Consequently, securities resembling economic catastrophe bonds should offer a large risk premium to compensate for their systematic risk.

Interestingly, we show that securities manufactured to resemble economic catastrophe bonds have relatively high prices, similar to single name securities with identical credit ratings. Credit ratings describe a security's expected payoffs in the form of its default likelihood and anticipated recovery value given default. However, because they contain no information about the state of the economy in which default occurs, they are insufficient for pricing. Nonetheless, in practice, many investors rely heavily upon credit ratings for pricing and risk assessment of fixed income securities, with large amounts of insurance and pension fund capital explicitly restricted to owning highly rated securities. In light of this behavior, the manufacturing of securities resembling economic catastrophe bonds emerges as the optimal mechanism for exploiting investors who rely on ratings for pricing. These securities will be the cheapest to deliver to investors demanding a given rating, but will trade at too high a price if valued based on rating-matched alternatives as opposed to proper risk-matched alternatives.

To study the risk properties of synthetic credit securities, we develop a simple state-contingent pricing framework. In the spirit of the Sharpe (1964) and Lintner (1965) CAPM, we use the realized market return as the relevant state space for asset pricing. This allows us to extract state prices from market index options using the technique of Breeden and Litzenberger (1978). Finally, to obtain state-contingent payoffs, we employ a modified version of the Merton (1974) structural credit model, in which asset values are driven by a common market factor. One of the well-documented weaknesses of structural models is that their reliance on lognormally-distributed asset values poses difficulty in pricing securities with low likelihoods of default. Because we use the structural approach solely to characterize default probabilities conditional on the level of the overall market, we only require *conditional* asset values to be lognormal, and therefore can remain agnostic about the distributional properties of the market return generating process. To price bonds and credit derivatives we then simply scale conditional payoffs by the option-implied state price density.

An attractive feature of this framework is that relying on the market state space preserves economic intuition throughout the pricing exercise, in contrast to popular statistics-heavy methods, which operate under a risk-neutral measure. The framework is assembled from classic insights on well developed markets, allowing the risks and prices of various securities to be consistently compared across markets.

Using the state price density extracted from index options, we calibrate the structural model to match the empirically observed credit yield spread, and then show that the replicating yield spread and the actual yield spread have similar dynamics, suggesting that the two markets are reasonably integrated. Our pricing model explains roughly 35% of the variation in weekly credit spread changes of a broad credit default swap index, which compares favorably to existing *ad hoc* specifications. At the same time, the market prices of highly-rated credit derivatives on this index are significantly higher than their risk-matched alternatives. In particular, we estimate that an investor who purchases the AAA-rated tranche of a collateralized debt obligation (CDO) 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 more 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 corporate bonds and bond portfolios to market index options 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 \bar{p}_D , the bond's price will satisfy,

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

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 priced 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, Ω . 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,

$$\bar{p}_D \equiv \int_{\omega \in \Omega} p_D(\omega) f(\omega) d\omega \quad (4)$$

where \bar{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

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

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, $f(\omega)$,

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

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, $1 - \bar{p}_D$, 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 $\underline{\Omega}$. Conversely, the least expensive asset pays off with certainty on a set of measure $1 - \bar{p}_D$, 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 $\bar{\Omega}$. 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 (1964) and 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 unconditional default probability of \bar{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 \bar{p}_D , and the sets of states for which they yield unit payoffs correspond to the previously identified $\underline{\Omega}$ and $\bar{\Omega}$. To price these options analytically it is convenient to specialize to

¹The 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.

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,

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

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 \bar{p}_D are given by,

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

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

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 \bar{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 \bar{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 *tranch*ed, 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 clientele 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 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

²Although 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 characteristics and pricing properties of a digital tranche carry over to the tranche structures traded in real-world credit markets.

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, \bar{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 limiting argument.³ We depart from previous implementations of the structural models in two respects. First, we assume asset returns satisfy a CAPM relationship, which allows us to derive state-contingent expectations of the tranche values for all realizations of the market return. The majority of our analysis is carried out conditional on the realization of the market return, allowing us to remain agnostic about the details of the market return generating process. When comparing to the existing literature, it will be helpful to make unconditional statements about default probabilities. For these comparisons, we will make the auxiliary assumption that log market returns are normally distributed. By allowing the firms' asset value processes to be correlated through the common market factor we are also able to capture their common exposure to macroeconomic conditions, and introduce default dependency.⁴ A similar approach is adopted in Hull, Predescu and White's (2006) Monte Carlo study of credit spreads and CDO tranche prices. Second, we value the state-contingent payoff expectations of bonds and CDO tranches by applying state prices extracted from long-dated index options. This ensures that we correctly capture the risk premia investors demand for assets which fail to pay off in states with high marginal utility, and allows us to raise the average predicted spread, without overstating the risks associated with volatility or leverage – a key challenge emphasized by Eom, Helwege and Huang (2004) in their

³See Eom, Helwege and Huang (2004) and references therein, for a comprehensive survey of the empirical performance of structural models. The authors find that the Merton (1974) model has a tendency to underestimate credit spreads when estimated model parameters are used.

⁴For an early implementation of a single-factor based model see Vasicek (1987, 1991). Schonbucher (2000) provides an overview of factor models for portfolio credit risk. Zhou (2001) examines the ability of structural models to capture default correlations through asset correlations.

survey of structural models.⁵

We begin with the assumption that firm asset values are characterized by the following stochastic differential equation,

$$\frac{dA_i}{A_i} = (r_f + \beta \cdot \eta)dt + \beta\sigma_m dZ_m + \sigma_\epsilon dZ_i, \quad (9)$$

where r_f is the riskless rate, β is the CAPM beta of the asset returns on the market portfolio, η is the (total) market risk premium and σ_m and σ_ϵ and the market and idiosyncratic asset return volatilities, respectively. While we require dZ_i to be a Gaussian diffusion, we allow dZ_m to follow an arbitrary mean-zero stochastic process. We make the common assumption that a firm defaults if the terminal value of its assets, A_T , falls below the face value of debt, D .⁶ 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,

$$p_D(r_{M,T}) = \Phi \left[\frac{\ln \frac{D}{A_t} - \left(r_f + \beta \left(\frac{r_{M,T}}{T} - r_f \right) - \frac{\sigma_\epsilon^2}{2} \right) \cdot T}{\sigma_\epsilon \sqrt{T}} \right], \quad (10)$$

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 $p_D(r_{M,T})$.

Under Merton's (1974) structural model the (percentage) loss given default is an endogenous variable determined by the shortfall between the terminal realization of the asset value and the face value of debt,

$$\tilde{L}_i(r_{M,T}) = \left(\frac{D - \tilde{A}_{i,T}(r_{M,T})}{D} \right) \cdot \mathbf{1}_{\tilde{A}_{i,T}(r_{M,T}) \leq D} \quad (11)$$

where we have defined a default indicator variable, $\mathbf{1}_{\tilde{A}_{i,T}(r_{M,T}) \leq D}$, that takes on a value of one, when a firm's terminal asset value falls below the face value of debt, D .⁷ To facilitate tractability, we adopt a somewhat more reduced form approach to modeling firm-specific losses, while maintaining the implications of the structural model with regard to the default process. Specifically, we assume that recovery rates in default are exogenous and independent of the firm's terminal asset value, $\tilde{A}_{i,T}$.

⁵Although Eom, Helwege and Huang (2004) conclude that empirical implementations of structural models produce rather imprecise estimates bond yield spreads, Schaefer and Strebulaev (2005) find that the comparative statics produced by structural models can be used to successfully hedge corporate bonds using equities.

⁶Black and Cox (1976) assume an alternative default process, in which default occurs at the first hitting time of the firm's asset value to a default threshold.

⁷An attractive feature of this modeling assumption is that low (high) realizations of the market return coincide with high (low) conditional default probabilities, and low (high) recovery rates, capturing the procyclical nature of recovery rates (Altman (2006)).

Consequently, while the terminal asset value continues to determine whether a firm has defaulted, its realization does not affect the recovery value.

Under this modified assumption, the conditional portfolio loss – given by an equal-weighted sum of the firm-specific losses, \tilde{L}_i – can be expressed as follows,

$$\tilde{L}_p(r_{M,T}) = \frac{1}{N} \sum_{i=1}^N \tilde{L}_i \cdot \mathbf{1}_{\tilde{A}_{i,T}(r_{M,T}) \leq D} \quad (12)$$

The percentage firm specific loss, \tilde{L}_i , is a random variable between $[0, 1]$, with mean l and variance v^2 . Although, we assume that the mean loss given default is independent of the realization of the common factor, $r_{M,T}$, procyclicality in recovery rates can be trivially incorporated by making the mean loss given default, l , a function of $r_{M,T}$. Under our assumption, the unconditional mean portfolio loss is identical under the objective and risk-neutral measures.⁸ To the extent that recovery rates covary positively with the realization of the market return, our assumption leads to a downward bias in the amount of systematic risk.

In what follows, we focus on the pricing of a digital tranche, which pays one dollar when the (percentage) portfolio loss is less than 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. We restrict our attention to these idealized tranches because the payoff to a non-infinitesimally tight tranche – i.e. a tranche with distinct lower and upper attachment points – can be replicated by a *strictly positive* combination of the digital tranches. Consequently, the pricing properties of the basis assets (i.e. the digital tranches) will be inherited by the composite asset.

In order to price the digital tranches, we first need to characterize their state-contingent payoffs, which requires an assessment of their conditional default probability,

$$p_D^X(r_{M,T}, N) = \text{Prob} \left(\tilde{L}_p(r_{M,T}) \leq X \right) \quad (13)$$

Unfortunately, closed-form expressions for the tranche default probability are elusive for moderate values of N .⁹ Only in the limit of a large homogenous portfolio, $N \rightarrow \infty$, can one derive an analytical expression for the tranche default probability (Vasicek (1987, 1997)). To do this note that the weak law of large numbers guarantees that the conditional portfolio loss converges to its

⁸Reduced form models employing the fractional recovery of market value convention fix the mean loss given default under the risk-neutral measure, L^Q (see Duffie and Singleton (2003)).

⁹A natural approach to this problem is to compute the characteristic function for the conditional portfolio loss, and then invert the Laplace transform to obtain the portfolio loss distribution function. However, the inverse transform is intractable for plausible firm-specific loss distributions. Alternative approaches involve copula-based simulations (Schonbucher (2002)).

mean in probability,

$$\begin{aligned} \lim_{N \rightarrow \infty} \tilde{L}_p(r_{M,T}) &= E \left[\frac{1}{N} \sum_{i=1}^N \tilde{L}_i \cdot \mathbf{1}_{\tilde{A}_{i,T}(r_{M,T}) \leq D} \right] \\ &= \frac{1}{N} \sum_{i=1}^N E \left[\tilde{L}_i \right] \cdot E \left[\mathbf{1}_{\tilde{A}_{i,T}(r_{M,T}) \leq D} \right] = l \cdot p_D(r_{M,T}) \quad \text{a.s.} \end{aligned} \quad (14)$$

Consequently, if we let $\hat{r}_{M,T}$ denote the value of the market return for which the portfolio loss converges to X , the conditional tranche default probability will be zero (one) when the realized market return, $r_{M,T}$, is above (below) $\hat{r}_{M,T}$,

$$\lim_{N \rightarrow \infty} p_D^X(r_{M,T}, N) = \mathbf{1}_{r_{M,T} \leq \hat{r}_{M,T}} \quad (15)$$

The corresponding, unconditional tranche default probability, which determines the tranche's credit rating, is given by,

$$\bar{p}_D^X = \int_{-\infty}^{\infty} \mathbf{1}_{r_{M,T} \leq \hat{r}_{M,T}} f(r_{M,T}) dr_{M,T} = F(\hat{r}_{M,T}) \quad (16)$$

where $f(\cdot)$ is the probability distribution function of the T -period market return. 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. If the continuously compounded market return is normal the strike price of the limiting digital market call option in moneyness space is given by $K^* = \exp(\hat{r}_{M,T})$, where,

$$\hat{r}_{M,T} = \left(r_f + \eta - \frac{\sigma_m^2}{2} \right) \cdot T - \sigma_m \sqrt{T} \cdot \Phi^{-1}(1 - \bar{p}_D^X) \quad (17)$$

We formalize the limiting pricing properties of the digital tranche in the following proposition.

Proposition 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 \bar{p}_D^X . As the number of securities in the portfolio underlying the tranche converges to infinity, $N \rightarrow \infty$, 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. When market returns are normal the price of the limiting call is given by,*

$$\lim_{N \rightarrow \infty} P_0^{X,N} = e^{-r_f T} \cdot \Phi \left[\Phi^{-1}(1 - \bar{p}_D^X) - \frac{\eta}{\sigma_m} \cdot \sqrt{T} \right] \quad (18)$$

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. In particular, 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 ($1 - p_D^X(r_{M,T}, N)$) with the state price ($q(r_{M,T})$) across all possible states,

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

The effect of increasing N on the tranche price, while holding its unconditional default probability constant, depends on how an increase in the number of underlying securities reallocates the probability of default from states with low marginal utility to states with high marginal utility. Specifically, because the realization of $r_{M,T}$ orders states in ascending order of marginal utility, if $\frac{\partial p_D^X(r_{M,T}, N)}{\partial N}$ is positive (negative) for low (high) market returns the tranche price will decline as N increases. Intuitively, the price of the digital tranche will 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 A.

Proposition 2 *If $\frac{\partial p_D^X(r_{M,T}, N)}{\partial N}$ is positive (negative) for realizations of the market return, $r_{M,T}$, below (above) $\hat{r}_{M,T}$, the price of a digital tranche with attachment point, X , written on a portfolio of N identical assets, $P_0^{X,N}$, will be monotonically decreasing in N , and will converge to the price of the limiting digital market call, $P_0^{X,\infty}$, as $N \rightarrow \infty$.*

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 September 2004 to September 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 implied volatility. 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 1 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 weekly correlations of each series in levels and Panel C reports correlations of first differences. Changes in long-term volatility are positively correlated with all bond spreads other than the AA and A. The CDX and all of the tranche spreads 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, as well as all bond spreads with a rating of BBB and lower.

3 Calibrating the Bond Pricing Model

Our calibration relies on the structural model to produce a state contingent payoff function for the CDX, which is then combined with an empirical estimate of the state-price density obtained from 5-year index options, to match the observed CDX yield. In other words, we project the payoffs of the CDX into the space of market returns using the structural model, and then use Arrow-Debreu prices to arrive at the CDX price. By requiring that our model match the CDX price on each day, we are able to calibrate a daily time series of the underlying parameters (leverage ratios, idiosyncratic asset volatility, and asset beta) for the representative firm in the credit default index. Our calibration effectively assumes that the CDX is comprised of bonds issued by N identical firms, so our estimated parameters are best thought of as characterizing the “average” firm in the index.

An implicit assumption of our calibration procedure, consistent with industry practice, is that the CDX spread reflects the risk-adjusted compensation for the expected loss given default, and is unaffected by tax or liquidity considerations. Longstaff, Mithal and Neis (2005) argue that a lack of supply constraints, the ease of entering and exiting credit default swap arrangements, and the contractual nature of the swaps, ensure that the market is less sensitive to liquidity and convenience yield effects, in contrast to the corporate bond market.

To verify the performance of our model we perform a variety of robustness checks. First, we

compare the performance of two parametric implied volatility functions used in constructing the state price density. Second, we show that our calibration procedure allows us to attain high R^2 in forecasting CDX yield changes at various frequencies. This ensures that the state-contingent replicating portfolio implied by the structural model shares the risk characteristics of the CDX index. We then show how the model can be used to price CDO tranches, as well as, construct simple replicating strategies involving put spreads on the market index.

3.1 Extracting the State Price Density

In order to extract the state price density, we exploit the fact that the prices of Arrow-Debreu securities can be recovered from option data. Given the market prices of European call options with maturity T and strike prices K , $C_t(K, T)$, Breeden and Litzenberger (1978) have shown that the price of an Arrow-Debreu security is equal to the second derivative of the call price function with respect to the strike price.¹⁰

$$q(x) = \left. \frac{\partial^2 C_t(K, T)}{\partial K^2} \right|_{K=M \cdot S} \quad (19)$$

where M is a moneyness level, defined as the ratio of the option strike price to the prevailing spot price. The formula for the Arrow-Debreu prices is particularly simple when the underlying follows a log-normal diffusion. However, as is now well established, index options exhibit a pronounced volatility smile, which suggests that deep out-of-the-money states are more expensive, than would be suggested by a simple log-normal diffusion model. To account for this, we derive the analog of the Breeden and Litzenberger (1978) result in the presence of a volatility smile. Specifically, we account for the fact that the Black-Scholes implied volatility is a function of the option strike price. Rewriting the call option price as $C_t(K, \sigma(K), T)$ and applying the chain rule we obtain,¹¹

$$q(x) = \left(\frac{\partial^2 C_t}{\partial K^2} + \frac{d\sigma}{dK} \cdot \left(\frac{\partial C_t}{\partial K \partial \sigma} + \frac{\partial^2 C_t}{\partial \sigma^2} \cdot \frac{d\sigma}{dK} \right) + \frac{\partial C_t}{\partial \sigma} \cdot \frac{d^2 \sigma}{dK^2} \right)_{K=M \cdot S} \quad (20)$$

If $\sigma_t(K, T)$ is given in closed-form, so are the prices of the Arrow-Debreu securities and the corresponding risk neutral density. As intuition suggests, the Arrow-Debreu prices now depend on the slope and curvature of the implied volatility smile, as well as the cross-partial effect of changes in the strike on option value.

To compute the Arrow-Debreu state prices, we fit an implied volatility function to the observed market option prices on each day by minimizing the pricing errors, and then substitute the function into (20). The 5-year implied volatilities are nearly linear in moneyness over the range for which we have observations (moneyness of 0.7 to 1.3). We choose two simple parametric forms for the implied volatility function, each of which is roughly linear around moneyness of 1.0, produces strictly

¹⁰To obtain the corresponding risk-neutral probabilities one simply multiplies the prices of the Arrow-Debreu securities by a factor of $e^{r(T-t)}$

¹¹This following expression is properly defined if and only if the implied volatility function, $\sigma_t(K, T)$, is twice differentiable in K .

positive implied volatilities, and is twice differentiable (see Appendix B for details). In particular, we assume that the implied volatility function is either exponential or hyperbolic tangent. Our specification allows us to compute all of the requisite derivatives in closed-form and is similar in spirit to the parametric methods employed by Rosenberg and Engle (2001) and Bliss and Panigirtzoglou (2004).¹²

Figure 2 displays the calibrated 5-year state prices and implied volatility functions as of each CDX initiation date. Both parametric forms produce average 5-year at-the-money implied volatilities of around 20% and about 10% at very high moneyness levels, but differ in the left tail of the moneyness distribution. The exponential implied volatility function averages nearly 40% at a moneyness of 0, while the hyperbolic tangent implied volatility function averages closer to 30% in this range. The implied state price densities tend to have very fat left tails between moneyness levels of 0 to 0.5 under both parametric assumptions, reflecting the high price of bad economic states expressed in the index options market.

3.2 Implying the Conditional Payoff

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 under the objective measure is given by the sum of payoffs on the N bonds underlying the index,

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

where \tilde{L}_i is the firm-specific loss given default, and R is the mean recovery rate, equal to one minus the expected loss given default, l . Since we have assumed that defaults and recovery values are independent of each other, the expectation of the product of the loss random variable and the default indicator is given by the product of their expectations. In our calibrations, we fix the mean recovery rate at 40%, consistent with industry practice.¹³ Moreover, because we assume losses are purely idiosyncratic and their conditional mean is independent of the realization of the market return, $r_{M,T}$, the mean recovery rates are identical under the objective and risk-neutral measures.

Finally, 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

¹²Ait-Sahalia and Lo (1998) propose an alternative, non-parametric method for extracting the state-price density, but their method requires large amounts of data and is not amenable to producing estimates at the daily frequency. For a literature review on methods for extracting the risk-neutral density from option prices see Jackwerth (1999) or Brunner and Hafner (2003).

¹³The Lehman Brothers bond-implied CDS spreads assume 40% recovery rates. Altman and Kishore (1996) and Duffie and Singleton (1999) report that the median recovery rate for senior unsecured bonds is roughly equal to 50%.

CDX payoff and the state price, $q(r_{M,T})$, across all realization of the market return we obtain the price of the CDX (or, equivalently, the price of the representative bond),

$$P_{CDX}(R) \equiv \int_{-\infty}^{\infty} E^P[CDX(r_{M,T})] \cdot q(r_{M,T}) dr_{M,T} \quad (22)$$

In pricing the CDX index, we exploit the fact that the state prices can be found in closed-form for our parametric specification of the implied volatility function. By using the state prices extracted from long-dated equity index options, we effectively 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 finds that the pricing of individual credit default swaps is consistent with the option-implied pricing kernel.

Using the above equation, combined with our empirical estimate of the state prices, we vary the underlying model parameters, $\{\frac{D}{A_t}, \beta_a, \sigma_\varepsilon\}$, until we match the the CDX price. In general, there may be multiple solutions to this non-linear equation since we only have one constraints and three parameters. Consequently, we also require that the model implied equity volatility or beta, match its empirical counterpart. Since the CDX is comprised of investment grade securities issued by large U.S. corporations, we typically require that the model implied equity beta equal one.

3.3 Evaluating the Model

The calibration procedure enables the model to match the CDX spread exactly at each point in time. However, assessing the model’s ability to accurately characterize the priced risks of corporate bonds requires that the model dynamics also match the dynamics of the CDX. To analyze the joint effectiveness of our model and calibration procedure at capturing the time series dynamics of the CDX, we regress weekly changes in CDX spreads on the change predicted by the model, as well as changes in the model’s underlying variables. Table 2 reports the output from these regressions. We calculate the model predicted change from time t to $t+1$ as the difference between the model yield at time $t+1$, using parameters calibrated at time t , and the actual yield at time t . The model predicted change is highly statistically significant with a large R^2 for both implementations of the model. The model predicted change has a t -statistic of 7.16 and an R^2 of 0.34 under the exponential implied volatility function and a t -statistic of 6.83 and an R^2 of 0.32 under the hyperbolic tangent implied volatility function.¹⁴ The change in the index level and the 5-year implied volatility are the most significant of the model’s variables in univariate and multiple regressions, but lose significance when the model predicted change is included. This suggests that the model has identified several relevant variables, and that the structure imposed by the model is helpful in explaining the dynamics of the CDX. In addition, the explanatory power of the model compares favorably to other empirical investigations into the determinants of credit spread changes for corporate bonds and CDSs (Collin-

¹⁴Results are essentially identical if we use a static put option on the market to match the market risks of the CDX, where the strike price and quantity are chosen to match the CDX yield in combination with a maturity-matched riskfree bond.

Dufresne, Goldstein, Martin (2001) and Zhang, Zhou, Zhu (2006)).

As a second exercise to assess the model's implications, we decompose the CDX spread into compensation for expected loss and risk. Following convention in the credit literature, we report the ratio of risk-neutral default intensity, λ^Q , to objective default intensity, λ^P . To compute the unconditional default probability under the historical measure, we make an auxiliary assumption that the terminal distribution of the market is lognormal. Then, using the calibrated parameters, we can compute the Merton model implied default probability, \bar{p}_D , and its corresponding objective default intensity, as,

$$\lambda^P = -\frac{1}{T} \ln \left(1 - \Phi \left[\frac{\ln \frac{D}{A_t} - \left(r_f + \beta\eta - \frac{\beta_a^2 \sigma_m^2 + \sigma_\epsilon^2}{2} \right) \cdot T}{\sqrt{\beta_a^2 \sigma_m^2 + \sigma_\epsilon^2} \cdot \sqrt{T}} \right] \right) \quad (23)$$

where the expression inside $\Phi(\cdot)$ is the distance-to-default. To get a sense of the quantity of systematic risk in the CDX we can compare the objective default intensity to its risk-neutral counterpart. 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 , through the following formula,

$$\begin{aligned} e^{-y_{CDX} \cdot T} &= e^{-r_f \cdot T} \cdot E^Q \left[1 \cdot \mathbf{1}_{\tilde{A}_{i,T} > D} + (1 - \tilde{L}_i) \cdot \mathbf{1}_{\tilde{A}_{i,T} \leq D} \right] \\ &= e^{-r_f \cdot T} \cdot \left(e^{-\lambda^Q \cdot T} + R \cdot (1 - e^{-\lambda^Q \cdot T}) \right) \end{aligned} \quad (24)$$

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). 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 ratio of λ^Q to λ^P reflects the relative importance of the risk-premium term (i.e. second term in (3)) in the pricing of a defaultable bond and 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 in states with high marginal utility the higher the value of the ratio. Elton et al. (2001) and Berndt et al. (2004) find evidence suggesting that corporate bond yield spreads contain important risk premia in addition to compensation for the expected default loss; Hull, Predescu and White (2005) report ratios of $\frac{\lambda^Q}{\lambda^P}$ that average 9.8 for A-rated bonds and 5.1 for BBB-rated bonds between 1996 and 2004. 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.

Historically, the representative firm included in the CDX index has had a credit rating of *BBB* or *A*. For example, Kakodar and Martin (2004) report that the CDX index had an average rating of BBB+ at the end of June 2004. Our calibration produces results consistent with this finding.

Figure 3 displays the daily time series of the calibrated objective default intensity, yield spread, and ratio of $\frac{\lambda^Q}{\lambda^P}$ for the CDX. The mean calibrated default intensity for the CDX is 20bps, corresponding to a 5-year default probability of 0.99%, which is between that for A-rated (0.50%) and BBB-rated (2.08%) bonds as reported in Hamilton et al. (2005). The ratio of $\frac{\lambda^Q}{\lambda^P}$ averages 3.8, ranging from 2.8 to 6.6.

4 Pricing Credit Derivatives

The Merton (1974) credit model integrated with the CAPM produces state-contingent payoffs for bonds and bond portfolios. These security-level payoffs are conditional on the realized market return, which allows for pricing via the market index option implied state price density. In other words, this pricing framework provides a direct link between the bond market and the index option market. The calibration procedure ensures consistency in price levels between the two markets, and results in similar price dynamics, suggesting that these two markets are reasonably integrated. The question now is whether the prices of tranches issued on the bond portfolio are consistent with their market risks.

This unified framework makes pricing derivatives simple. Having recovered the time series of model parameters (asset beta, leverage level, and idiosyncratic volatility) of the representative bond in the CDX, we can simulate state-contingent payoffs for the CDX. This requires one additional assumption about the conditional distribution of the firm-level loss. We assume the percentage loss given default for each issue comes from a beta distribution with mean of 60% (1-recovery rate) and standard deviation of 10%. The terms of each derivative security (i.e. tranche) define its payoff as a function of the underlying security's payoff. In this case, the tranche payoff is defined as a call spread on CDX losses, or equivalently, a put spread on the CDX payoff. The state-contingent tranche payoff is identified by applying the contract terms to each simulated outcome, and pricing is completed, as before, using the Arrow-Debreu prices. The state-contingent payoffs for the CDX and one of its senior tranches are displayed in Figure 4.

Table 3 presents a comparison of the spreads predicted by the models with the spreads offered by each of the CDX tranches. In particular, for both implementations of the model, we report the time series mean of the actual and model spreads, the correlation between weekly yields and changes in yields, the 5-year model implied default probability, and the mean ratio of $\frac{\lambda^Q}{\lambda^P}$. The credit risk ratio, $\frac{\lambda^Q}{\lambda^P}$, for the tranches is calculated by dividing the annualized model yield spread by the loss rate (i.e. the annualized yield spread that would be obtained by discounting at the riskfree rate).

Across all tranches, our model predicts significantly greater spreads than are present in the data. 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 four times as large as in the data. Figures 5-8 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 5-8 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 six.

On the other hand, correlations in weekly spread levels 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 model and its corresponding CDX tranche are driven by common economic risk factors.

The model implied 5-year default probabilities for the 7-10 tranche appear to be slightly higher than the historical average for single-name AAA-rated bonds. The 10-15 tranche is perhaps more in line with the historical default probabilities for AAA-rated single name bonds. Interestingly, both the 7-10 and 10-15 tranches have large credit risk ratios (18.2 and 59.3, respectively) relative to the historical average for single name AAA-rated bonds (16.8 according to Hull, Predescu, and White (2005)). This is consistent with the earlier prediction that a much larger portion of the tranche yield spread represents compensation for market risk, as opposed to expected losses, than is the case for single-name bonds with similar default probabilities.

4.1 A Short-Cut for Derivatives on Large Portfolios

When the number of issues in the CDX becomes large, the (conditional) CDX payoff converges in probability to its (conditional) mean. Similarly, the tranche payoff, which can be thought of as a call spread on the portfolio loss, converges in probability to a payoff resembling a put spread on the market. We can 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,

$$E[L_p(r_{M,T})] = (1 - R) \cdot p_D(r_{M,T}) \tag{25}$$

Setting this value equal to X and solving for $\exp(r_{M,T})$, yields the corresponding put price strike, K_X ,¹⁵

$$K_X = \exp \left\{ \frac{1}{\beta} \left(\ln \frac{D}{A_t} - \left(r_f \cdot (1 - \beta) - \frac{\sigma_\epsilon^2}{2} \right) \cdot T - \sigma_\epsilon \sqrt{T} \cdot \Phi^{-1} \left(\frac{X}{1 - R} \right) \right) \right\} \tag{26}$$

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 .

¹⁵All option strike prices are expressed as a fraction of the spot price. Note, this computation assumes that continuously compounded market returns are normally distributed.

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 are 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 4), and remains constant over the life of the tranche. The yield of the replicating portfolio is simply $-\frac{1}{T} \ln V_t$.

Panel B of Table 3 compares the actual tranche yields to the market-risk-matched put spreads. The results using the put-spread approximation are very similar to those using the simulated tranche payoff functions. The replicating portfolios offer considerably larger yield spreads than the CDO tranches, the correlations between weekly changes of the actual and model spreads are relatively high, and the ratio of $\frac{\lambda^Q}{\lambda^P}$ increases dramatically with the seniority of the tranche. Finally, it is useful to note that the put-spread approximation represents a static replicating portfolio, making it attractive from an implementation perspective.

4.2 Robustness

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 written on the S&P 500 index with different upper and lower strike prices. Panel A presents yields calculated using an exponential implied volatility function. Panel B presents yields calculated with a hyperbolic tangent implied volatility function.

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 0.30 and a lower strike price of 0.25 or an upper strike price of 0.35 and a lower strike price of 0.20. 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 65 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 tranche traded at a yield spread of 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 only a the 30-0 put spread (simply short a put) 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 over 70% 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). It is worth noting that since 1871, the US

stock market has never declined more than 60% over any 5-year period, and outside the 1930s, its maximum 5-year decline is 36%.¹⁶

As an additional robustness check, we also conduct the analysis in Table 4 using the longest-dated S&P options trading on the CBOE. At the end of our sample (September 11, 2006), the CBOE traded put options with a maximum maturity of December 2008. The annual yields offered by put spreads with this maturity were similar to the OTC options used above. For instance, a put spread with moneyness attachment points of 0.54 and 0.46 offered an 86 basis point yield spread, one with attachment points of 0.62 and 0.54 offered 171 basis points, and one with attachment points of 0.69 and 0.62 offered 307bps. In Panel C, the comparable yield spreads offered by the five-year options are roughly 138, 186, and 247bps, respectively.

A final concern with our approach—and with structural credit models in general—is that we rely on the equivalence between market and firm-specific returns in terms of their relation to debt values. That is, our analysis implicitly assumes that a 10 percent decline in a firm’s equity value has the same impact on its default likelihood regardless of whether this decline is firm-specific or market-wide. This assumption may be problematic if factors such as sentiment or discount rate news have a strong influence on overall valuations but say nothing about a specific firm’s cash flows and therefore its ability to repay its obligations. Of course, there are also reasons why market returns may be more informative than excess returns about a firm’s default likelihood. For instance, if delevering (e.g. selling assets) is more difficult when many firms are under financial pressure, a market-driven decline in firm value may portend more poorly for bondholders than an equivalent, firm-specific decline.

Figure 9 presents some evidence on this issue by plotting the percentage of publicly-traded firms that saw their bonds downgraded each month against the past one-year return on the market. The plot shows a strong negative relationship (correlation = -0.60) has held between the two series since 1987. This suggests that, at least over this limited recent period, past market returns are important in explaining a given firm’s downgrade likelihood. To examine the relative importance of market returns, we estimate a logistic regression that uses lagged one-year firm returns separated into market and excess returns to explain monthly firm-level downgrade events. Lagged market returns enter the regression with a coefficient that is at least as large and significant as that of the excess firm return. Thus, our implicit assumption that market-wide and firm-specific returns have a similar relation to a firm’s credit quality does not appear inconsistent with the U.S. experience over the past 19 years.

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

¹⁶Since 1801, the maximum 5-year decline experienced by the FTSE is 55% and since 1914, the maximum 5-year decline of the Nikkei is 54%.

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 ($N \rightarrow \infty$), 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, single-name counterparts, 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 state-contingent approach for the valuation of fixed income securities, which has the virtue of preserving economic intuition even when applied to complex derivatives. Our pricing strategy for collateralized debt obligation tranches is to identify packages of other investable securities that deliver identical payoffs conditional on the market return. 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 assume the economic risks inherent in senior CDO tranches can, with equivalent economic exposure, earn roughly three times more compensation by writing out-of-the-money put spreads on the market. We argue that this discrepancy has much to do with the fact that credit rating agencies are willing to certify senior CDO tranches as “safe” when, from an asset pricing perspective, they are quite the opposite.

References

- [1] Ait-Sahalia, Yacine and Andrew W. Lo, 1998, Nonparametric estimation of state-price densities implicit in financial asset prices, *Journal of Finance* 53, 499-547.
- [2] Altman, Edward, 2006, Default Recovery rates and LGD in Credit Risk Modeling and Practice: An Updated Review of the Literature and Empirical Evidence, NYU working paper.
- [3] Altman, Edward and Vellore, 1996, Almost Everything You Wanted to Know About Recoveries on Defaulted Bonds, *Financial Analysts Journal* 52, 57-64.
- [4] Berndt, Antje, Rohan Douglas, Darrell Duffie, Mark Ferguson and David Schranz, 2005, Measuring Default-Risk Premia from Default Swap Rates and EDFs, Stanford GSB working paper.
- [5] Black, Fischer and John C. Cox, 1976, Valuing Corporate Securities: Some Effects of Bond Indenture Provisions, *Journal of Finance* 31, 351-367.
- [6] Black, Fischer and Myron Scholes, 1973, The pricing of options and corporate liabilities, *Journal of Political Economy* 81, 617-654.
- [7] Bliss, Robert and Nikolaos Panigirtzoglou, 2004, Option-Implied Risk Aversion Estimates, *Journal of Finance* 59, 407-446.
- [8] Breeden, Douglas, and Robert Litzenberger, 1978, Prices of state-contingent claims implicit in option prices, *Journal of Business*, 621-651.
- [9] Brunner, Bernhard and Reinhold Hafner, 2003, Arbitrage-free Estimation of the Risk-neutral Density from the Implied Volatility Smile, *Journal of Computational Finance* 7, 75-106.
- [10] Collin-Dufresne, Pierre, Robert S. Goldstein, and J. Spencer Martin, 2001, The Determinants of Credit Spread Changes, *Journal of Finance* 56, 2177-2208.
- [11] 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).
- [12] Delianedis, Gordon, and Robert Geske, 2001, The Components of Corporate Credit Spreads: Default, Recovery, Tax, Jumps, Liquidity, and Market Factors, UCLA working paper.
- [13] DeMarzo, Peter, 2005, The pooling and tranching of securities: A model of informed intermediation, *Review of Financial Studies*, 1-35.
- [14] Duffie, Darrell, and Kenneth J. Singleton, 1999, Modeling Term Structures of Defaultable Bonds, *Review of Financial Studies* 12, 687-720.
- [15] Duffie, Darrell, and Kenneth Singleton, 2003, *Credit Risk*, Princeton University Press.

- [16] Elton, Edwin J., Martin J. Gruber, Deepak Agrawal, and Christopher Mann, 2001, Explaining the Rate Spread on Corporate Bonds, *Journal of Finance* 56, 247-277.
- [17] Eom, Young Ho, Jean Helwege, and Jing-Zhi Huang, 2004, Structural Models of Corporate Bond Pricing, *Review of Financial Studies*, 17, p. 499-544.
- [18] Huang, Jing-zhi, and Ming Huang, 2003, How Much of the Corporate-Treasury Yield Spread is Due to Credit Risk?, Penn State University working paper.
- [19] Hull, John, Mirela Predescu, and Alan White, 2005, Bond prices, default probabilities, and risk premiums, *Journal of Credit Risk*, 53-60.
- [20] Hull, John, Mirela Predescu, and Alan White, 2006, The Valuation of Correlation-dependent Credit Derivatives Using a Structural Model, University of Toronto working paper.
- [21] Jackwerth, Jens C., Option Implied Risk-Neutral Distributions and Implied Binomial Trees: A Literature Review, *Journal of Derivatives* 7, 66-82.
- [22] Jones, E. Philip, Scott P. Mason, and Eric Rosenfeld, 1984, Contingent Claims Analysis of Corporate Capital Structures: an Empirical Investigation, *Journal of Finance* 39, 611-625
- [23] Kakodkar, Atish and Barnaby Martin, 2004, The Standardized Tranche Market Has Come Out of Hiding, *Journal of Structured Finance*, Fall, 76-81.
- [24] 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.
- [25] Longstaff, Francis, Sanjay Mithal and Eric Neis, 2005, Corporate Yield Spreads: Default Risk or Liquidity? New Evidence from the Credit Default Swap Market, *Journal of Finance* 60, 2213-2253.
- [26] Longstaff, Francis, and Arvind Rajan, 2007, An Empirical Analysis of the Pricing of Collateralized Debt Obligations, *Journal of Finance*, forthcoming.
- [27] Merton, Robert C., 1973, Theory of rational option pricing, *Bell Journal of Economics and Management Science*, 141-183.
- [28] Merton, Robert C., 1974, On the pricing of corporate debt: The risk structure of interest rates, *Journal of Finance*, 449-470.
- [29] Pan, Jun, and Kenneth Singleton, 2006, Default and recovery implicit in the term structure of sovereign CDS spreads, MIT working paper.
- [30] Rosenberg, Joshua and Robert F. Engle, 2002, Empirical pricing kernels, *Journal of Financial Economics*, 64, 341-372.

- [31] Schaefer, Stephen and Ilya Strebulaev, 2005, Structural Models of Credit Risk are Useful: Evidence from Hedge Ratios on Corporate Bonds, Stanford GSB working paper.
- [32] Schonbucher, Philipp, 2000, Factor Models for Portfolio Credit Risk, Bonn University working paper.
- [33] Sharpe, William F., 1964, Capital asset prices: A theory of market equilibrium under conditions of risk, *Journal of Finance* 19, 425-442.
- [34] Vasicek, Oldrich, 1987, Probability of Loss on Loan Portfolio, KMV Corporation.
- [35] Vasicek, Oldrich, 1991, Limiting Loan Loss Distribution, KMV Corporation.
- [36] Zhang, Benjamin, Hao Zhou, and Haibin Zhu, 2006, Explaining Credit Default Spreads with the Equity Volatility and Jump Risks of Individual Firms, BIS working paper.
- [37] Zhou, Chunsheng, 2001, An Analysis of Default Correlations and Multiple Defaults, *Review of Financial Studies* 14, 555-576.

A Tranche Price Convergence as $N \rightarrow \infty$

We are interested in establishing the direction and monotonicity of the convergence of the tranche price as $N \rightarrow \infty$, when the unconditional tranche default probability is held fixed at \bar{p}_D^X . Because the unconditional default probability is held constant we have,

$$\frac{\partial}{\partial N} \int_{-\infty}^{\infty} p_D^X(r_{M,T}, N) \cdot f(r_{M,T}) dr_{M,T} = 0 \quad (\text{A1})$$

where $f(r_{M,T})$ is the probability density function of the T -period market return. Since increasing N has no effect on the unconditional tranche default probability, it also has no effect on the unconditional expectation of the tranche payoff. 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 Arrow-Debreu expression for the tranche price with respect to N ,

$$\frac{\partial P_0^{X,N}}{\partial N} = - \int_{-\infty}^{\infty} \frac{\partial p_D^X(r_{M,T}, N)}{\partial N} \cdot q(r_{M,T}) dr_{M,T} \quad (\text{A2})$$

Multiplying and dividing the state price, $q(r_{M,T})$, by the probability density function of the T -period market return, $f(r_{M,T})$, we obtain,

$$\frac{\partial P_0^{X,N}}{\partial N} = - \int_{-\infty}^{\infty} \frac{\partial p_D^X(r_{M,T}, N)}{\partial N} \cdot \frac{q(r_{M,T})}{f(r_{M,T})} \cdot f(r_{M,T}) dr_{M,T} \quad (\text{A3})$$

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 sign the above derivative we require that the partial derivative of the tranche default probability with respect to the number of underlyings, N , satisfy the following condition,

$$\begin{cases} \frac{\partial p_D^X(r_{M,T}, N)}{\partial N} > 0 & \text{for } r_{M,T} < \hat{r}_{M,T} \\ \frac{\partial p_D^X(r_{M,T}, N)}{\partial N} < 0 & \text{for } r_{M,T} > \hat{r}_{M,T} \end{cases} \quad (\text{A4})$$

where $\hat{r}_{M,T}$ is the critical value of the market return at which the probability of default switches from one to zero for the large, homogenous tranche ($N \rightarrow \infty$). This condition allows us to split the integral in (A3) at $\hat{r}_{M,T}$ and apply the mean value theorem.

Although we cannot show that the above property holds for a general value of N since we do not have an expression for $p_D^X(r_{M,T}, N)$, it is possible to show that a one-sided Chebyshev bound on the tranche default probability satisfies condition (A4). Consequently, we take the above condition as given, apply the mean-value theorem twice and exploit the monotonicity of the stochastic discount factor to sign the effect of increasing N on the tranche price, while holding the unconditional default probability constant,

$$\begin{aligned} \frac{\partial P_0^{X,N}}{\partial N} &= - \left(\int_{-\infty}^{\hat{r}_{M,T}} \frac{\partial p_D^X(r_{M,T}, N)}{\partial N} \cdot m(r_{M,T}) \cdot f(r_{M,T}) dr_{M,T} + \int_{\hat{r}_{M,T}}^{\infty} \frac{\partial p_D^X(r_{M,T}, N)}{\partial N} \cdot m(r_{M,T}) \cdot f(r_{M,T}) dr_{M,T} \right) \\ &= - \left(\bar{m} \cdot \int_{-\infty}^{\hat{r}_{M,T}} \frac{\partial p_D^X(r_{M,T}, X, N)}{\partial N} \cdot f(r_{M,T}) dr_{M,T} + \underline{m} \cdot \int_{\hat{r}_{M,T}}^{\infty} \frac{\partial p_D^X(r_{M,T}, N)}{\partial N} \cdot f(r_{M,T}) dr_{M,T} \right) \end{aligned} \quad (\text{A5})$$

where \bar{m} is some number in $[m(-\infty), m(\hat{r}_{M,T})]$, and \underline{m} is some number in $[m(\hat{r}_{M,T}), m(\infty)]$. However, because $m(r_{M,T})$ is monotonically declining in the market return we immediately have $\bar{m} > \underline{m}$. 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,

$$\int_{-\infty}^{\hat{r}_{M,T}} \frac{\partial p_D(r_{M,T}, X, N)}{\partial N} \cdot f(r_{M,T}) dr_{M,T} + \int_{\hat{r}_{M,T}}^{\infty} \frac{\partial p_D(r_{M,T}, X, N)}{\partial N} \cdot f(r_{M,T}) dr_{M,T} = 0 \quad (\text{A6})$$

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

$$\frac{\partial P_0^{X,N}}{\partial N} = -(\bar{m} \cdot \xi - \underline{m} \cdot \xi) = \xi \cdot (\underline{m} - \bar{m}) < 0 \quad (\text{A7})$$

which indicates that the price of the digital tranche is monotonically declining in N .

B Fitting the Implied Volatility Smile

We consider two parametric forms for the implied volatility function when fitting the prices of five-year S&P 500 index options:

$$\sigma_t(x, T) = a + b \cdot \exp(-c \cdot x) \quad (\text{B1})$$

$$\sigma_t(x, T) = a + b \cdot \tanh(-c \cdot \ln x) \quad (\text{B2})$$

where x is the option moneyness, defined as the ratio of the option strike price to the prevailing spot price. These functional forms for the volatility skew have a variety of attractive features. First, they can generate an approximately linear skew for options whose strike prices are close to at-the-money, matching the stylized facts for long-dated options. Second, the functions are bounded above and below, allowing us to control the magnitude of the implied volatility outside of the domain of strike prices for which we observe option prices. This approach is more elegant than assuming a constant implied volatility outside the range of observable option prices (Shimko (1993), Campa et al. (1997)), since it avoids problems of non-differentiability and ensures that prices are martingales. The solution proposed by Brown and Toft (1999) is most similar to ours. They assume that the implied volatility function converges to some exogenously specified value. Instead, we “fit” this asymptotic value by requiring that the implied risk-neutral density allow us to consistently price a riskless bond and a futures contract (i.e. an option with a strike of $K = 0$). Lastly, the simple parametric forms we chose facilitate the analytical computation of the derivatives appearing in the formula for the Arrow-Debreu prices.

For example, for the exponential implied volatility function, the first and second derivatives of the implied volatility function with respect to the strike price (or, equivalently, moneyness) are given by:

$$\frac{d\sigma_t(x, T)}{dx} = -bc \cdot \exp(-c \cdot x) \quad (\text{B3})$$

$$\frac{d^2\sigma_t(x, T)}{dx^2} = bc^2 \cdot \exp(-c \cdot x) \quad (\text{B4})$$

The corresponding derivatives for the hyperbolic tangent implied volatility function are:

$$\frac{d\sigma_t(x, T)}{dx} = \left(\frac{b \cdot c}{x}\right) \cdot \operatorname{sech}^2(c \cdot \ln x) \quad (\text{B5})$$

$$\frac{d^2\sigma_t(x, T)}{dx^2} = -\left(\frac{b \cdot c}{x^2}\right) \cdot \operatorname{sech}^3(c \cdot \ln x) \cdot (\cosh(c \cdot \ln x) + 2c \cdot \sinh(c \cdot \ln x)) \quad (\text{B6})$$

The remainder of the partial derivatives appearing in formula for the Arrow-Debreu state prices can be obtained by differentiating the Black-Scholes formula with respect to its various parameters. In particular, we obtain:

$$\left. \frac{\partial^2 C_t}{\partial K^2} \right|_x = e^{-rT} \cdot \frac{\phi(d_2(x))}{x\sigma_t(x, T)\sqrt{T}} \quad (\text{B7})$$

$$\left. \frac{\partial C_t}{\partial \sigma} \right|_x = \phi(d_1(x)) \cdot \sqrt{T} = e^{-rT} \cdot x \cdot \phi(d_2(x)) \cdot \sqrt{T} \quad (\text{B8})$$

$$\left. \frac{\partial^2 C_t}{\partial \sigma^2} \right|_x = \frac{d_1(x) \cdot d_2(x)}{\sigma_t(x, T)} \cdot \phi(d_1(x)) \cdot \sqrt{T} = e^{-rT} \cdot \frac{x \cdot d_1(x) \cdot d_2(x)}{\sigma_t(x, T)} \cdot \phi(d_2(x)) \cdot \sqrt{T} \quad (\text{B9})$$

$$\left. \frac{\partial C_t}{\partial K \partial \sigma} \right|_x = \frac{2d_1(x) \cdot \phi(d_1(x))}{x \cdot \sigma_t(x, T)} = e^{-rT} \cdot \frac{2d_1(x) \cdot \phi(d_2(x))}{\sigma_t(x, T)} \quad (\text{B10})$$

where:

$$d_1(x) = \frac{\ln \frac{1}{x} + (r_f + \frac{1}{2} \cdot \sigma_t(x, T)^2) \cdot T}{\sigma_t(x, T) \cdot \sqrt{T}} \quad (\text{B11})$$

$$d_2(x) = d_1 - \sigma_t(x, T) \cdot \sqrt{T} \quad (\text{B12})$$

Substituting these expressions back into the formula for the Arrow-Debreu prices, yields a closed-form representation for the state prices implied by our parametrization of the implied volatility function.

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

This table reports summary statistics of various credit market securities. The 5-year bond-implied CDS spreads by rating group are provided by Lehman Brothers. The CDX series is the Dow Jones CDX North America Investment Grade index of 5-year 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 5yVol. The skew is the linear slope between 5-year at-the-money volatility and 5-year 30% out-of-the-money volatility. The yield on 5-year constant maturity Treasury bonds is denoted Rf.

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

| | VIX | 5yVOL | Skew | Rf | AA | A | BBB | BB | B | CDX | [3-7] | [7-10] | [10-15] | [15-30] |
|------|------|-------|-------|-----|------|------|------|-------|-------|------|-------|--------|---------|---------|
| Mean | 13.2 | 17.8 | -12.8 | 4.2 | 15.3 | 25.6 | 46.9 | 143.1 | 254.5 | 47.8 | 146.2 | 44.0 | 19.8 | 9.2 |
| Std | 1.8 | 1.3 | 0.8 | 0.5 | 2.2 | 2.6 | 5.5 | 22.9 | 36.8 | 6.9 | 56.5 | 24.8 | 9.0 | 3.7 |

Panel B: Correlations between weekly series

| | VIX | 5yVOL | Skew | Rf | AA | A | BBB | BB | B | CDX | [3-7] | [7-10] | [10-15] | [15-30] |
|---------|-------|-------|-------|-------|-------|------|------|-------|------|------|-------|--------|---------|---------|
| VIX | 1.00 | | | | | | | | | | | | | |
| 5yVOL | 0.38 | 1.00 | | | | | | | | | | | | |
| Skew | 0.23 | 0.00 | 1.00 | | | | | | | | | | | |
| Rf | 0.17 | 0.58 | 0.25 | 1.00 | | | | | | | | | | |
| AA | -0.02 | 0.55 | 0.05 | 0.40 | 1.00 | | | | | | | | | |
| A | -0.11 | 0.14 | -0.15 | -0.38 | 0.27 | 1.00 | | | | | | | | |
| BBB | 0.04 | 0.52 | -0.11 | -0.14 | 0.54 | 0.85 | 1.00 | | | | | | | |
| BB | 0.24 | 0.63 | -0.40 | -0.01 | 0.28 | 0.57 | 0.73 | 1.00 | | | | | | |
| B | -0.06 | 0.39 | -0.50 | -0.23 | 0.24 | 0.72 | 0.76 | 0.85 | 1.00 | | | | | |
| CDX | 0.07 | -0.14 | -0.53 | -0.63 | -0.22 | 0.52 | 0.39 | 0.52 | 0.64 | 1.00 | | | | |
| [3-7] | 0.18 | -0.43 | -0.05 | -0.83 | -0.48 | 0.32 | 0.13 | 0.08 | 0.20 | 0.66 | 1.00 | | | |
| [7-10] | 0.12 | -0.49 | 0.05 | -0.86 | -0.44 | 0.27 | 0.10 | -0.08 | 0.05 | 0.53 | 0.96 | 1.00 | | |
| [10-15] | 0.10 | -0.46 | -0.12 | -0.89 | -0.43 | 0.39 | 0.19 | 0.09 | 0.22 | 0.72 | 0.96 | 0.96 | 1.00 | |
| [15-30] | -0.03 | -0.32 | -0.39 | -0.66 | -0.39 | 0.51 | 0.26 | 0.33 | 0.46 | 0.89 | 0.69 | 0.60 | 0.77 | 1.00 |

Panel C: Correlations between changes in weekly series

| | VIX | 5yVOL | Skew | Rf | AA | A | BBB | BB | B | CDX | [3-7] | [7-10] | [10-15] | [15-30] |
|---------|-------|-------|-------|-------|-------|-------|------|------|------|------|-------|--------|---------|---------|
| VIX | 1.00 | | | | | | | | | | | | | |
| 5yVOL | 0.46 | 1.00 | | | | | | | | | | | | |
| Skew | 0.03 | -0.09 | 1.00 | | | | | | | | | | | |
| Rf | -0.09 | 0.06 | -0.07 | 1.00 | | | | | | | | | | |
| AA | -0.01 | -0.05 | -0.04 | -0.20 | 1.00 | | | | | | | | | |
| A | -0.15 | -0.05 | 0.04 | -0.26 | 0.71 | 1.00 | | | | | | | | |
| BBB | -0.07 | 0.11 | -0.02 | -0.36 | 0.63 | 0.85 | 1.00 | | | | | | | |
| BB | 0.27 | 0.18 | -0.06 | -0.57 | 0.24 | 0.33 | 0.47 | 1.00 | | | | | | |
| B | 0.25 | 0.22 | 0.05 | -0.48 | 0.22 | 0.34 | 0.45 | 0.77 | 1.00 | | | | | |
| CDX | 0.38 | 0.26 | -0.03 | -0.13 | -0.02 | -0.01 | 0.10 | 0.47 | 0.37 | 1.00 | | | | |
| [3-7] | 0.12 | 0.19 | -0.15 | 0.00 | 0.17 | 0.21 | 0.27 | 0.48 | 0.43 | 0.49 | 1.00 | | | |
| [7-10] | 0.09 | 0.12 | -0.15 | -0.02 | 0.17 | 0.17 | 0.22 | 0.30 | 0.28 | 0.46 | 0.84 | 1.00 | | |
| [10-15] | 0.16 | 0.15 | -0.15 | -0.10 | 0.18 | 0.15 | 0.25 | 0.39 | 0.31 | 0.58 | 0.80 | 0.82 | 1.00 | |
| [15-30] | 0.13 | 0.16 | -0.10 | -0.02 | 0.02 | 0.08 | 0.12 | 0.35 | 0.28 | 0.58 | 0.65 | 0.62 | 0.77 | 1.00 |

Table 2
Regressions Explaining the Time Series of Changes in CDX Yield Spreads (9/2004 – 9/2006)

The dependent variable is the weekly change in the CDX yield spread. The CDX series is the Dow Jones CDX North America Investment Grade index of 5-year credit default swaps. The model predicted change in the CDX yield spread, $E[\Delta Y_{CDX,Exp}]$ and $E[\Delta Y_{CDX,Tanh}]$, is calculated as the difference between the model yield at time $t+1$, using parameters calibrated at time t , and the actual yield at time t , using the exponential and hyperbolic tangent implied volatility state price densities, respectively. The yield on 5-year constant maturity Treasury bonds is denoted R_f . S&P refers to the level of the S&P 500 index. Five-year at-the-money implied volatility is denoted as 5yVol. The skew is the linear slope between 5-year at-the-money volatility and 5-year 30% out-of-the-money volatility. The adjusted R-square is denoted R^2 , t -statistics are in parenthesis, and the number of observations are in square brackets.

| Intercept | $E[\Delta Y_{CDX,Exp}]$ | $E[\Delta Y_{CDX,Tanh}]$ | ΔR_f | $\Delta S\&P$ | $\Delta 5yVol$ | $\Delta Skew$ | R^2 [N] |
|-------------------|-------------------------|--------------------------|--------------------|--------------------|--------------------|--------------------|----------------|
| 0.0000 (-0.29) | 0.2702 (7.19) | | | | | | 0.3455 [97] |
| 0.0000 (-0.28) | | 0.2445 (6.87) | | | | | 0.3249 [97] |
| 0.0000 (-0.31) | | | -0.0555 (-2.05) | | | | 0.0324 [97] |
| 0.0000 (-0.03) | | | | -0.0092 (-5.46) | | | 0.2309 [97] |
| 0.0000 (-0.84) | | | | | 0.0285 (4.78) | | 0.1852 [97] |
| 0.0000 (-0.62) | | | | | | 0.0529 (1.26) | 0.0060 [97] |
| 0.0000 (-0.13) | | | -0.0275 (-1.16) | -0.0068 (-3.60) | 0.0147 (2.23) | 0.0537 (1.49) | 0.2858 [97] |
| 0.0000 (-0.44) | 0.5430 (2.39) | | 0.0215 (0.68) | 0.0071 (1.24) | -0.0134 (-1.08) | -0.0462 (-0.92) | 0.3291 [97] |
| 0.0000 (-0.35) | | 0.3916 (1.76) | 0.0080 (0.25) | 0.0044 (0.71) | -0.0080 (-0.61) | 0.0285 (0.85) | 0.3106 [97] |

Table 3
Comparison of Actual and Model Implied Tranche Spreads

The actual weekly CDX tranche spreads corresponds to various tranches of the Dow Jones CDX North America Investment Grade index of 5-year credit default swaps. The model yield is calculated a Merton (1974) / CAPM credit model calibration procedure. The credit risk ratio, λ^Q/λ^P , for the tranches is calculated by dividing the model yield spread by the loss rate (i.e. the yield spread that would be obtained by discounting at the riskfree rate). The mean 5-year default probability is calculated using the time series average of the weekly objective default intensity.

| Panel A: Model assumes continuum of index option strikes | | | | | | |
|---|-------------------------|--------------------------|---------------------------------|---|-------------------------------------|------------------------------|
| Tranche | Mean Model Spread [bps] | Mean Actual Spread [bps] | Correlation of Model and Actual | Correlation of Model and Actual (Changes) | Mean 5-Year Default Probability (%) | Mean λ^Q / λ^P |
| <i>Exponential Implied Volatility Function</i> | | | | | | |
| 3%-7% | 184 | 147 | 0.78 | 0.44 | 4.37 | 5.3 |
| 7%-10% | 110 | 44 | 0.40 | 0.39 | 0.57 | 18.2 |
| 10%-15% | 78 | 20 | 0.42 | 0.51 | 0.17 | 59.3 |
| 15%-30% | 44 | 9 | 0.50 | 0.37 | 0.03 | 576.5 |
| <i>Hyperbolic Tangent Implied Volatility Function</i> | | | | | | |
| 3%-7% | 215 | 147 | 0.58 | 0.54 | 5.55 | 4.6 |
| 7%-10% | 124 | 44 | 0.22 | 0.44 | 0.76 | 15.0 |
| 10%-15% | 80 | 20 | 0.40 | 0.51 | 0.23 | 44.6 |
| 15%-30% | 34 | 9 | 0.71 | 0.35 | 0.04 | 330.8 |
| Panel B: Model assumes a static index put spread is available at each CDX initiation date | | | | | | |
| Tranche | Mean Model Spread [bps] | Mean Actual Spread [bps] | Correlation of Model and Actual | Correlation of Model and Actual (Changes) | Mean 5-Year Default Probability (%) | Mean λ^Q / λ^P |
| <i>Exponential Implied Volatility Function</i> | | | | | | |
| 3%-7% | 200 | 147 | 0.19 | 0.56 | 3.72 | 5.6 |
| 7%-10% | 115 | 44 | 0.22 | 0.32 | 0.41 | 22.2 |
| 10%-15% | 82 | 20 | 0.31 | 0.49 | 0.15 | 72.0 |
| 15%-30% | 47 | 9 | 0.30 | 0.28 | 0.03 | 576.5 |
| <i>Hyperbolic Tangent Implied Volatility Function</i> | | | | | | |
| 3%-7% | 234 | 147 | 0.20 | 0.53 | 4.50 | 5.0 |
| 7%-10% | 130 | 44 | 0.26 | 0.31 | 0.52 | 18.4 |
| 10%-15% | 85 | 20 | 0.33 | 0.50 | 0.19 | 53.1 |
| 15%-30% | 37 | 9 | 0.25 | 0.20 | 0.04 | 308.5 |

Table 4
Yields on Various S&P 500 Put Spreads

The yield on the 5-year S&P 500 put spread is calculated 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 reports the average daily yield for various put spreads assuming an exponential implied volatility function. Panel B reports the average daily yield for various put spreads assuming a hyperbolic tangent implied volatility function. Panel C reports the daily yield for various put spreads on September 11, 2006, assuming an exponential implied volatility function. Panel D reports the daily yield for various put spreads on September 11, 2006, assuming a hyperbolic tangent implied volatility function.

Average Daily Yield from 9/2004 to 9/2006

| Panel A: Exponential Implied Volatility Function | | | | | | | | | | | |
|--|------|--------------------|------|------|------|------|------|------|------|------|------|
| | | Upper Strike Price | | | | | | | | | |
| | | 0.30 | 0.35 | 0.40 | 0.45 | 0.50 | 0.55 | 0.60 | 0.65 | 0.70 | 0.75 |
| Lower Strike Price | 0.00 | 17 | 22 | 28 | 34 | 42 | 50 | 59 | 68 | 80 | 92 |
| | 0.05 | 20 | 26 | 32 | 39 | 46 | 55 | 64 | 74 | 86 | 99 |
| | 0.10 | 24 | 30 | 37 | 44 | 52 | 61 | 70 | 81 | 93 | 106 |
| | 0.15 | 29 | 36 | 43 | 50 | 58 | 67 | 77 | 88 | 101 | 115 |
| | 0.20 | 35 | 42 | 49 | 56 | 65 | 74 | 85 | 96 | 109 | 124 |
| | 0.25 | 41 | 48 | 55 | 63 | 72 | 82 | 93 | 105 | 118 | 133 |
| | 0.30 | | 55 | 62 | 71 | 80 | 90 | 101 | 114 | 128 | 144 |
| | 0.35 | | | 70 | 79 | 88 | 99 | 111 | 124 | 139 | 155 |
| | 0.40 | | | | 88 | 98 | 109 | 121 | 135 | 150 | 168 |
| | 0.45 | | | | | 108 | 119 | 132 | 147 | 163 | 181 |
| | 0.50 | | | | | | 131 | 145 | 160 | 177 | 196 |
| | 0.55 | | | | | | | 159 | 175 | 193 | 213 |
| | 0.60 | | | | | | | | 191 | 210 | 232 |
| 0.65 | | | | | | | | | 230 | 252 | |
| 0.70 | | | | | | | | | | 275 | |

| Panel B: Hyperbolic Tangent Implied Volatility Function | | | | | | | | | | | |
|---|------|--------------------|------|------|------|------|------|------|------|------|------|
| | | Upper Strike Price | | | | | | | | | |
| | | 0.30 | 0.35 | 0.40 | 0.45 | 0.50 | 0.55 | 0.60 | 0.65 | 0.70 | 0.75 |
| Lower Strike Price | 0.00 | 8 | 13 | 19 | 27 | 35 | 44 | 55 | 66 | 78 | 91 |
| | 0.05 | 10 | 15 | 22 | 30 | 39 | 49 | 60 | 71 | 84 | 98 |
| | 0.10 | 12 | 18 | 26 | 34 | 44 | 54 | 66 | 78 | 91 | 106 |
| | 0.15 | 16 | 23 | 31 | 40 | 50 | 61 | 73 | 86 | 99 | 115 |
| | 0.20 | 21 | 28 | 37 | 47 | 57 | 69 | 81 | 95 | 109 | 125 |
| | 0.25 | 27 | 35 | 45 | 55 | 66 | 78 | 91 | 105 | 120 | 136 |
| | 0.30 | | 44 | 54 | 64 | 76 | 89 | 102 | 116 | 132 | 149 |
| | 0.35 | | | 64 | 75 | 87 | 100 | 114 | 129 | 145 | 162 |
| | 0.40 | | | | 86 | 99 | 112 | 127 | 142 | 158 | 177 |
| | 0.45 | | | | | 112 | 125 | 140 | 156 | 173 | 192 |
| | 0.50 | | | | | | 139 | 155 | 171 | 189 | 208 |
| | 0.55 | | | | | | | 170 | 187 | 206 | 226 |
| | 0.60 | | | | | | | | 204 | 224 | 245 |
| 0.65 | | | | | | | | | 243 | 266 | |
| 0.70 | | | | | | | | | | 288 | |

Table 4 (Continued)

Yield on 9/11/2006

Panel C: Exponential Implied Volatility Function

| | | Upper Strike Price | | | | | | | | | |
|--------------------|------|--------------------|------|------|------|------|------|------|------|------|------|
| | | 0.30 | 0.35 | 0.40 | 0.45 | 0.50 | 0.55 | 0.60 | 0.65 | 0.70 | 0.75 |
| Lower Strike Price | 0.00 | 15 | 21 | 28 | 35 | 44 | 54 | 64 | 76 | 90 | 104 |
| | 0.05 | 18 | 24 | 32 | 40 | 49 | 59 | 70 | 83 | 97 | 112 |
| | 0.10 | 22 | 29 | 37 | 45 | 55 | 66 | 77 | 90 | 105 | 121 |
| | 0.15 | 27 | 35 | 43 | 52 | 62 | 73 | 85 | 99 | 114 | 130 |
| | 0.20 | 33 | 41 | 50 | 60 | 70 | 82 | 94 | 108 | 124 | 141 |
| | 0.25 | 41 | 49 | 58 | 68 | 79 | 91 | 104 | 119 | 135 | 153 |
| | 0.30 | | 57 | 67 | 77 | 89 | 101 | 115 | 130 | 147 | 166 |
| | 0.35 | | | 76 | 87 | 99 | 112 | 127 | 143 | 160 | 180 |
| | 0.40 | | | | 98 | 111 | 125 | 140 | 156 | 175 | 195 |
| | 0.45 | | | | | 123 | 138 | 154 | 171 | 190 | 212 |
| | 0.50 | | | | | | 152 | 169 | 187 | 208 | 230 |
| | 0.55 | | | | | | | 186 | 205 | 226 | 249 |
| | 0.60 | | | | | | | | 224 | 247 | 271 |
| | 0.65 | | | | | | | | | 269 | 295 |
| 0.70 | | | | | | | | | | 321 | |

Panel D: Hyperbolic Tangent Implied Volatility Function

| | | Upper Strike Price | | | | | | | | | |
|--------------------|------|--------------------|------|------|------|------|------|------|------|------|------|
| | | 0.30 | 0.35 | 0.40 | 0.45 | 0.50 | 0.55 | 0.60 | 0.65 | 0.70 | 0.75 |
| Lower Strike Price | 0.00 | 7 | 13 | 20 | 28 | 38 | 49 | 61 | 74 | 88 | 104 |
| | 0.05 | 9 | 15 | 23 | 32 | 42 | 54 | 67 | 80 | 95 | 111 |
| | 0.10 | 11 | 18 | 27 | 37 | 48 | 60 | 73 | 88 | 103 | 120 |
| | 0.15 | 14 | 22 | 32 | 43 | 55 | 68 | 82 | 97 | 113 | 130 |
| | 0.20 | 20 | 28 | 39 | 50 | 63 | 77 | 92 | 107 | 124 | 142 |
| | 0.25 | 27 | 36 | 48 | 60 | 73 | 88 | 103 | 119 | 137 | 156 |
| | 0.30 | | 46 | 58 | 71 | 85 | 100 | 116 | 133 | 151 | 170 |
| | 0.35 | | | 70 | 84 | 98 | 114 | 130 | 148 | 166 | 186 |
| | 0.40 | | | | 98 | 113 | 129 | 146 | 164 | 183 | 204 |
| | 0.45 | | | | | 128 | 144 | 162 | 180 | 200 | 222 |
| | 0.50 | | | | | | 161 | 179 | 198 | 219 | 241 |
| | 0.55 | | | | | | | 197 | 217 | 238 | 261 |
| | 0.60 | | | | | | | | 237 | 259 | 283 |
| | 0.65 | | | | | | | | | 282 | 307 |
| 0.70 | | | | | | | | | | 332 | |

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 |

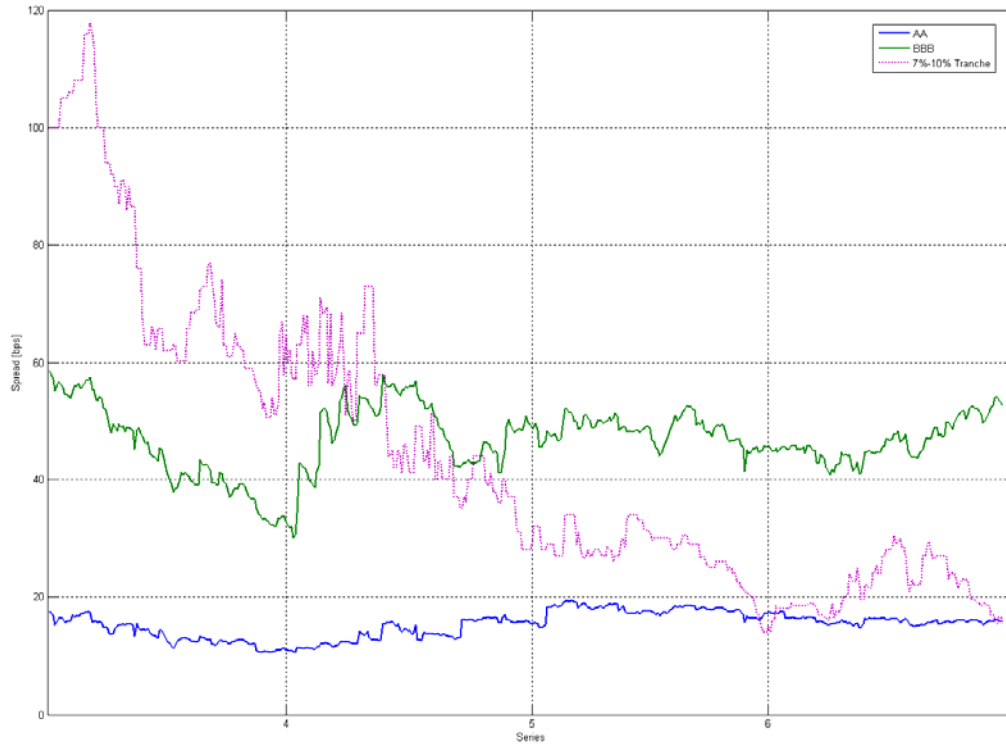


Figure 1. Time series of the 7%-10% tranche spread and credit spreads for bond indices (9/2004 – 9/2006). The AA and BBB 5-year bond-implied CDS spreads are provided by Lehman Brothers. The CDX tranche spread corresponds to the 7%-10% tranche of the Dow Jones CDX North America Investment Grade index of 5-year credit default swaps.

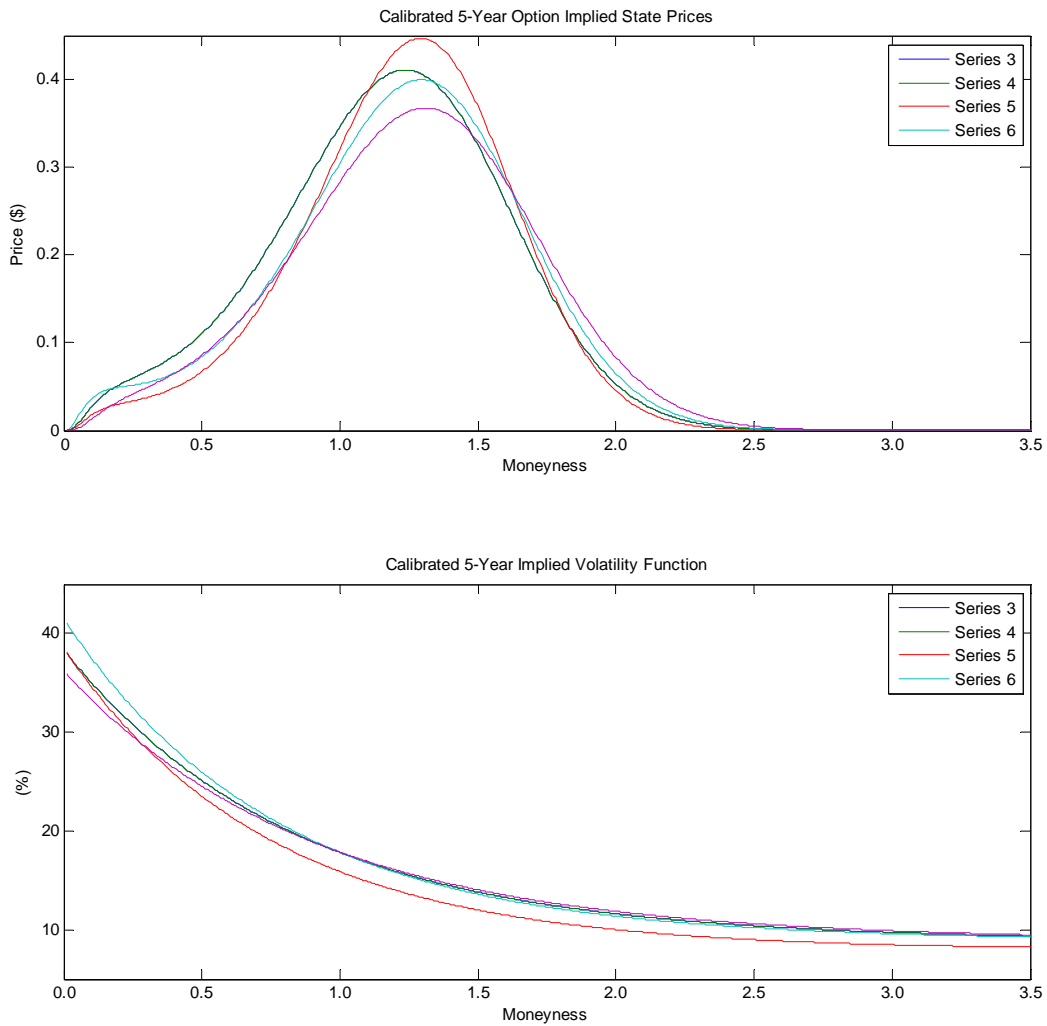


Figure 2. Calibrated 5-year state prices and implied volatility functions as of CDX initiation dates.

Panel A: Exponential implied volatility function.

The top panel displays state prices calculated using the technique of Breeden and Litzenberger (1978), adjusted to account for the implied volatility skew. The bottom panel shows the parametric implied volatility function fitted using prices of 5-year S&P 500 index options. The observed option prices come from securities with normalized strike prices ranging from 0.70 to 1.30, at 0.05 increments.

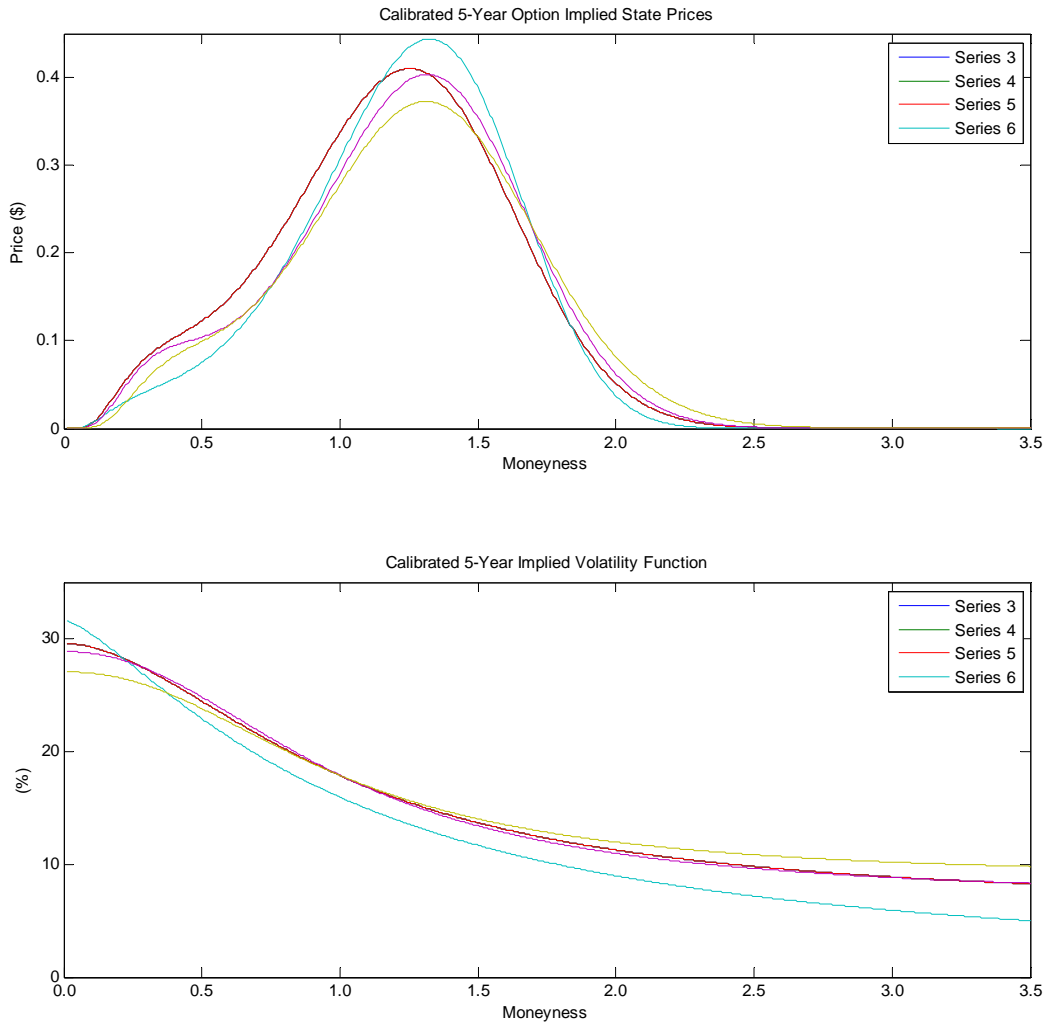


Figure 2. Calibrated 5-year state prices and implied volatility functions as of CDX initiation dates.

Panel B: Hyperbolic tangent implied volatility function.

The top panel displays state prices calculated using the technique of Breeden and Litzenberger (1978), adjusted to account for the implied volatility skew. The bottom panel shows the parametric implied volatility function fitted using prices of 5-year S&P 500 index options. The observed option prices come from securities with normalized strike prices ranging from 0.70 to 1.30, at 0.05 increments.




Figure 3. Time series of calibrated CDX default intensity and spread.

The CDX yield, y_{CDX} , is given by the Dow Jones CDX North America Investment Grade index of 5-year credit default swaps. Each day, the Merton (1974) structural credit model is applied to 125 identical CAPM assets to determine the CDX default rate conditional on the market return. The conditional payoff for the CDX is calculated using a recovery rate, R , of 40%. The conditional payoff is priced with a market index option implied state price density, constructed from a fitted exponential implied volatility function. To calculate the objective default intensity, λ^P , we assume the terminal value of the market index is lognormal. The risk-neutral default intensity, λ^Q , is calculated as $y_{CDX}/(1-R)$.

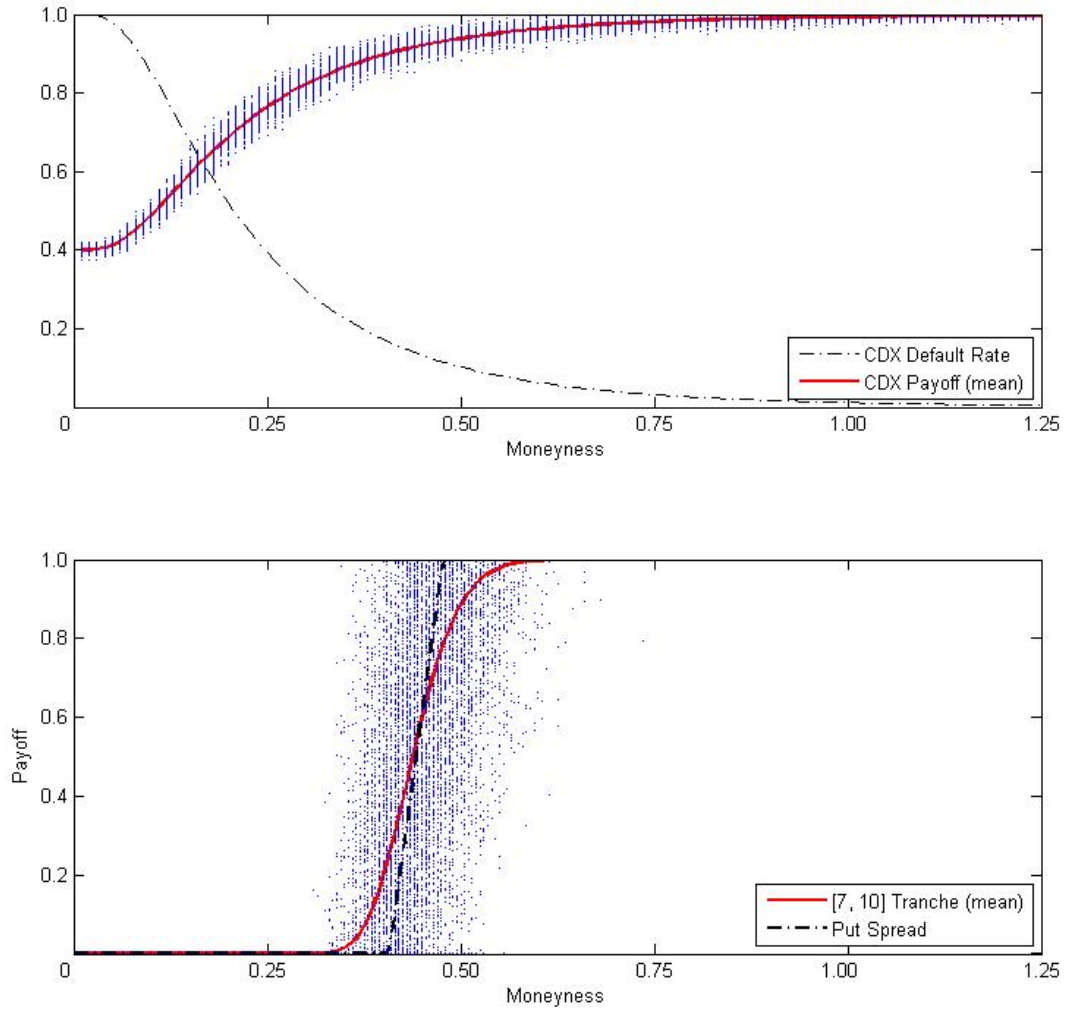


Figure 4. State Contingent CDX default rate, CDX payoff, and the 7%-10% tranche payoff.

The top panel of the figure plots the calibrated firm default rate as a function of the 5-year market return. The default rate characterizes the representative bond in the credit default index (CDX). The figure includes a sample of the CDX payoff realizations (blue dots) and their state contingent mean (red line). The payoffs are obtained via Monte Carlo simulation using calibrated firm parameters as of the last date in the sample. The mean recovery rate conditional on default is fixed at 40% of par. The bottom panel plots the corresponding payoffs to the 7%-10% tranche (blue dots) and their state-contingent mean (red line). The dashed black line depicts the payoff to an approximating put spread portfolio.



Figure 5. Time series of actual and model tranche spread (9/2004 – 9/2006).

The actual CDX tranche spread corresponds to the 3%-7% tranche on the Dow Jones CDX North America Investment Grade index of 5-year credit default swaps. The model yield is estimated by first applying the Merton (1974) credit model to 125 identical CAPM assets to determine the CDX default rate conditional on the market return. The conditional payoff for the CDX is calculated using a recovery rate, R , of 40%. The conditional payoff is priced with a market index option implied state price density, constructed from a fitted exponential implied volatility function. The model parameters are calibrated such that the model CDX yield matches the observed CDX yield. The conditional payoff for the tranche is simulated using the calibrated CDX parameters and priced using the state price density.



Figure 6. Time series of actual and model tranche spread (9/2004 – 9/2006).

The actual CDX tranche spread corresponds to the 7%-10% tranche on the Dow Jones CDX North America Investment Grade index of 5-year credit default swaps. The model yield is estimated by first applying the Merton (1974) credit model to 125 identical CAPM assets to determine the CDX default rate conditional on the market return. The conditional payoff for the CDX is calculated using a recovery rate, R , of 40%. The conditional payoff is priced with a market index option implied state price density, constructed from a fitted exponential implied volatility function. The model parameters are calibrated such that the model CDX yield matches the observed CDX yield. The conditional payoff for the tranche is simulated using the calibrated CDX parameters and priced using the state price density.



Figure 7. Time series of actual and model tranche spread (9/2004 – 9/2006).

The actual CDX tranche spread corresponds to the 10%-15% tranche on the Dow Jones CDX North America Investment Grade index of 5-year credit default swaps. The model yield is estimated by first applying the Merton (1974) credit model to 125 identical CAPM assets to determine the CDX default rate conditional on the market return. The conditional payoff for the CDX is calculated using a recovery rate, R , of 40%. The conditional payoff is priced with a market index option implied state price density, constructed from a fitted exponential implied volatility function. The model parameters are calibrated such that the model CDX yield matches the observed CDX yield. The conditional payoff for the tranche is simulated using the calibrated CDX parameters and priced using the state price density.



Figure 8. Time series of actual and model tranche spread (9/2004 – 9/2006).

The actual CDX tranche spread corresponds to the 15%-30% tranche on the Dow Jones CDX North America Investment Grade index of 5-year credit default swaps. The model yield is estimated by first applying the Merton (1974) credit model to 125 identical CAPM assets to determine the CDX default rate conditional on the market return. The conditional payoff for the CDX is calculated using a recovery rate, R , of 40%. The conditional payoff is priced with a market index option implied state price density, constructed from a fitted exponential implied volatility function. The model parameters are calibrated such that the model CDX yield matches the observed CDX yield. The conditional payoff for the tranche is simulated using the calibrated CDX parameters and priced using the state price density.




Figure 9. U.S. Bond Downgrade Rate and Lagged Market Returns (12/1987 – 4/2007).

Downgrades correspond to the fraction of publicly-listed US firms with S&P-rated debt that experienced a downgrade during the calendar month. The lagged one-year return corresponds to the cumulative return on the S&P 500 during the past 12 month period. Using individual firm returns, the corresponding logistic regression is estimated as follows:

$$\text{Prob}(\text{Downgrade}_{i,t}=1) = \frac{\exp X' \beta}{1 + \exp X' \beta}, \text{ where } X' \beta = b_0 + b_1 \cdot R_{m,t} + b_2 \cdot (R_{i,t} - R_{m,t}).$$

The estimated coefficients are displayed in the table below, with p -values reported in parenthesis.

| b_0 | b_1 | b_2 | N |
|--------|--------|--------|---------|
| 3.90 | 3.27 | 3.01 | 129,330 |
| (0.03) | (0.13) | (0.06) | |