

Economic Catastrophe Bonds

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

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 occurs. Such investors are likely to be attracted to securities whose payoffs resemble 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.

First draft: March 2007

This draft: October 2008

*Coval: Harvard Business School; *jcoval@hbs.edu*. Jurek: Bendheim Center for Finance, Princeton University; *jjurek@princeton.edu*. Stafford: Harvard Business School; *estafford@hbs.edu*. We thank John Campbell, Mikhail Chernov, John Cochrane, Pierre Collin-Dufresne, Francis Longstaff, Bob Merton, André Perold, Chester Spatt, and Jeremy Stein for valuable comments and discussions, and seminar participants at Boston College, Boston University, Georgia State University, Harvard Management Company, Harvard University, Ohio State University, Princeton University, MIT, NYU, University of Chicago, University of Pennsylvania, Virginia Tech, Yale University, the 2007 NBER Asset Pricing Summer Institute, the 2007 UNC/Duke Asset Pricing Conference, the 2008 Utah Winter Finance Conference, the 15th Mitsui Life Symposium on Financial Markets, and the 2008 Western Finance Association Conference. We are especially grateful to Citigroup for providing index options data.

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 and garner 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 Kenneth J. Arrow (1964) and Gerard 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 find that securities manufactured to resemble economic catastrophe bonds have relatively high prices, similar to those of 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 credit sensitive 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 supply 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 manufactured credit securities, we develop a simple state-contingent pricing framework. In the spirit of the William F. Sharpe (1964) and John 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 Douglas Breeden and Robert Litzenberger (1978). Index options are essentially bets on whether the index value at maturity will be above or below a prespecified level. When appropriately combined, they can be used to price a bet that pays one if the index closes at a given value and zero otherwise, and thus represent the Arrow-Debreu state price for that particular market realization. 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. This allows us to compute the payoffs of the underlying asset pool as a function of the realized market return. Since the tranches are derivatives written on the pool, their state-contingent payoffs can be determined simply by ap-

plying their contractual terms to the payoffs of the pool. Finally, to price the asset pool and its derivatives, we scale the mean state-contingent payoffs by the corresponding option-implied state prices. 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. 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.

One of the well-documented weaknesses of structural credit 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 realization of the market return, we only require the *conditional* distribution of asset values to be lognormal, while remaining agnostic about the distribution of the market factor. This seemingly minor modification has a significant impact. Using the state price density extracted from index options, we are easily able to calibrate the structural model to match the empirically observed credit yield spread for a broad credit default swap index (CDX). The resulting parameters are intuitive and stable through time. We also show that the replicating yield spread and the actual yield spread have similar dynamics, suggesting that the index options and corporate bond markets are reasonably integrated. Specifically, the pricing model explains roughly 30% of the variation in weekly credit spread changes of the CDX, which compares favorably to existing *ad hoc* specifications.

Remarkably, at the same time, we find that the market prices of senior CDX tranches are significantly higher than their risk-matched alternatives. We estimate that an investor who purchases a AAA-rated tranche of the CDX could earn a yield spread four to five times larger by bearing comparable economic risks in the index options market. This disparity survives a variety of model specifications and, importantly, is not driven by extremely poor economic states that are outside of the US historical experience. In particular, senior tranches remain significantly overpriced even when the possibility of a five-year market index decline of greater than 60% is ruled out. Over the period from 2004 through 2007, spreads of senior CDX tranches are highly similar to those of equivalently rated single name corporate bonds, suggesting that investors indeed viewed these to be the proper benchmark even though their underlying economic risks are highly dissimilar.

1 The Impact of Pooling and Tranching on Asset Prices

Structured finance activities proceed in two steps. In the first step, a number of similar securities (bonds, loans, credit default swaps, etc.) are *pooled* in a special purpose vehicle. In the second step, the cash flows of this portfolio are redistributed, or *trunched*, across a series of derivative securities. The absolute priority observed in redistributing cash flows among the derivative claims, called tranches, allows some of them to have a lower likelihood of default (expected loss) than the average security in the underlying portfolio. In turn, senior tranches are able to obtain a credit rating higher than the average credit rating of the securities in the reference portfolio.

This process is essentially the same as that undertaken when a corporation issues bonds of varying seniority along with equity. As a result of the prioritization of the claims issued against the firm’s asset pool, defaults of senior claims are less likely than those of the junior claims. Moreover, the prioritization of the claims causes the defaults of more senior claims to, on average, be associated with progressively worse economic outcomes. This is reflected by high ratios of yield spreads to expected loss rates for highly rated corporate bonds (Edwin J. Elton et al. (2001), Joost Driessen (2003), John Hull, Mirela Predescu, and Alan White (2005)).

The tranching of portfolios composed of securities that already have a tendency to concentrate risks in bad economic states, further concentrates these risks. We show that losses on the most senior tranches referencing an index of investment grade credit default swaps are largely confined to the worst economic states, 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.

In order to examine the impact of pooling and tranching on asset prices, the next section introduces a conceptual model of credit securities in the spirit of Arrow (1964) and Debreu (1959). This model is then adapted in Section 2, into a form that is suitable for calibration to market data, allowing us to explore the pricing of structured finance securities in actual credit markets.

1.1 The Economic Setting

In order to explore the pricing of derivatives referencing pools of risky bonds, common in structured finance, we first need to pre-specify the set of possible economic outcomes. Since risk-averse investors will be willing to pay more for securities that pay off in states where marginal utility of one dollar is high, it is natural to order the economic states, s , in order of declining marginal utility. In what follows, we assume that s entirely captures the set of states in the economy that are relevant to investors, providing an indexing of states from most adverse (recession) to most favorable (boom). In this economic setting – first introduced by Arrow (1964) and Debreu (1959) – the challenge of pricing assets boils down to specifying their state-contingent payoffs.

Since a typical risky bond is issued by a firm, whose cash flows – and consequently, ability to repay – are positively related to economic outcomes, it is also natural to assume that risky bonds are *economic assets*, and their payoffs are positively exposed to the economy’s outcomes.¹ To keep matters simple, we assume that at maturity, τ periods into the future, each bond will either pay one, or with some probability, $p_D(s)$, default and pay zero. The economic nature of the risky bonds is captured by allowing the individual securities’ default probabilities to be related to the economic state, s . In particular, we assume that each asset’s default risk is a declining function of s , such that the bonds are more likely to default in recessions than in booms ($\frac{\partial p_D(s)}{\partial s} < 0$). Finally, we denote the probability of a given state s occurring by $\pi(s)$.

In practice, the default risks of credit sensitive securities – risky bonds and tranches alike – are

¹Empirically, this shows up in the form of positive slope coefficients when bond returns are regressed on the returns to a broad equity index. See, for example, Eugene F. Fama and Kenneth R. French (1993).

frequently assessed by government-certified credit rating agencies. These agencies (e.g. Moody’s, S&P, Fitch, etc.) assign alphanumeric ratings to securities on the basis of their unconditional default probabilities or expected losses.² The attention paid by investors to credit ratings often compels issuers to target a desired rating and therefore issue claims with a specific level of default risk. Consequently, the comparative statics of our pricing analysis will typically hold the credit rating – as proxied by the unconditional default probability – of the securities under consideration fixed.

1.2 Characterizing Tranche Prices

To develop intuition for how structured finance activities impact the state-contingent payoffs and prices of the derivative claims, we begin by considering a tranche referencing an asset pool with a notional value of \$1, consisting of N homogeneous and equally-weighted risky bonds. The simplest possible tranche, called a digital tranche, is one that offers a payoff of one if the portfolio payoff exceeds the value of $1 - X$ and zero otherwise. The value X is referred to as a loss attachment point, since it specifies the maximal percentage loss on the underlying portfolio for which the tranche continues to make its contractual payments.³ It can be thought of as the amount of protection or “overcollateralization” provided to the claim by the prioritized payout structure.

If the defaults in the underlying asset portfolio are independent given the realization of the economic state, s , the number of defaulted securities in the underlying portfolio in a given state will have a binomial distribution, with parameter $p_D(s)$ and N trials. In practice, the underlying asset pools are quite large and the state-contingent default probabilities on the individual assets sufficiently high, such that the binomial distribution will be well approximated by a normal distribution with mean, $N \cdot p_D(s)$, and variance, $N \cdot p_D(s) \cdot (1 - p_D(s))$.⁴ This distributional approximation has the added convenience that it allows us to treat N as a continuous variable, and simplifies a number of the ensuing derivations without any loss of intuition. The ensuing results will hold under more general conditions, and in the next section we formally relax this assumption to ensure that it does not quantitatively influence our empirical results. Using this approximation, the state-contingent probability of observing the tranche default, $p_D^X(s)$, given by the probability that the losses on the underlying portfolio exceed the attachment point, X , is

$$p_D^X(s) = 1 - \Phi \left(\sqrt{N} \cdot \frac{(X - p_D(s))}{\sqrt{p_D(s) \cdot (1 - p_D(s))}} \right), \tag{1}$$

where $\Phi(\cdot)$ is the cumulative normal distribution. Figure 1 illustrates the state-contingent tranche

²In the simplified setting in this section, the recovery value in default is set to zero and therefore the unconditional default probability and expected loss coincide.

³In real-world credit market settings, to which we turn in the next section, tranches are defined by a lower (X) and upper (Y) attachment point, such that the tranche pays one if the portfolio loses less than X , zero if it loses more than Y , and a linearly scaled payoff if the portfolio loss is between X and Y .

⁴This approximation can also be interpreted as saying that within a given state, s , the percentage portfolio losses are normally distributed with mean, $p_D(s)$, and variance, $\frac{p_D(s) \cdot (1 - p_D(s))}{N}$. Although the normal distribution counterfactually allows for negative losses, the likelihood of such draws shrinks dramatically with N .

payoff, given by $1 - p_D^X(s)$, and shows that – holding X fixed – increasing the number of securities in the underlying portfolio shifts payoffs from states with high marginal utility, to states with low marginal utility. However, as N increases, the unconditional tranche default probability will also vary, causing the tranche’s credit rating to change.

The unconditional tranche default probability can be computed from:

$$p_D^X = \int_s p_D^X(s) \cdot \pi(s) ds. \tag{2}$$

From this, it is easy to see that the claim’s default probability is declining in the amount of overcollateralization, X . Less obvious is how the default likelihood is impacted by N . In particular, for states $s < p_D^{-1}(X)$, the state-contingent probability of observing a tranche fail is increasing in N , while for states $s > p_D^{-1}(X)$, it is decreasing in N . The net effect of a change in the number of underlying securities on the unconditional tranche default probability, p_D^X , will therefore depend not only on the magnitudes of the derivatives in the two regions of the state space, but also on the relative likelihoods of observing economic outcomes in these two regions. Intuitively, for senior tranches – characterized by large values of X – the likelihood of observing economic states in which default occurs is likely to be sufficiently low, such that the second effect dominates, and an increase in N results in a decrease in the unconditional probability of default, p_D^X .

More central to our purposes is the question of how increasing the size of the asset pool, N – a proxy for the level of diversification – changes the price of a tranche, when the unconditional default probability is fixed at some level p^* , i.e. when the tranche’s credit rating is held fixed. Using the state-contingent pricing method of Arrow (1964) and Debreu (1959), if the value of one unit of wealth in state s is given by $q(s)$ – the Arrow-Debreu price for state s – the price of any asset can be obtained by summing the expected state-contingent payoffs multiplied by their corresponding state prices. For example, the price of each τ -period risky bond in the underlying portfolio can be obtained from

$$B_\tau = \int_s (1 - p_D(s)) \cdot q(s) ds. \tag{3}$$

A derivative claim, or tranche, against an underlying asset pool that is structured to achieve a certain level of default risk can be priced in an equivalent manner. In particular, if we define $X(p^*, N)$ as the loss attachment point that produces a digital tranche with default risk, p_D^X equal to p^* , when the underlying portfolio is comprised of N bonds, the price of the claim with τ periods to maturity can be computed from

$$B_\tau^{X(p^*, N)} = \int_s \Phi \left(\sqrt{N} \cdot \frac{(X(p^*, N) - p_D(s))}{\sqrt{p_D(s) \cdot (1 - p_D(s))}} \right) \cdot q(s) ds. \tag{4}$$

By analogy to the bond example, the first term in the integral represents the expected state-contingent tranche payoff and is equal to one minus the state-contingent tranche default probability, $p_D^{X(p^*, N)}(s)$; and the second term represents the corresponding state price. By varying the tranche

attachment point, X , one can form an entire sequence of tranches, indexed by N , whose unconditional default probability is held fixed at p^* . This attachment point is denoted by $X(p^*, N)$. We show that the values of the tranches in this sequence have the following property.

Proposition 1 *Holding the unconditional default probability constant at p^* , the value of a digital tranche, $B_\tau^{X(p^*, N)}$, declines as the number of securities in the underlying bond portfolio, N , increases.*

Proof: See Appendix A.

The intuition underlying Proposition 1 is that increasing N , while targeting a fixed level of default risk, reallocates payoffs from states with high marginal utility (recessions) to states with low marginal utility (booms). Consequently, despite having an unaltered credit rating, the securities in the sequence offer progressively less protection against economic catastrophe. In order to bear this increased level of systematic risk, the marginal investor demands additional compensation, causing the tranche price to fall.

1.3 Cheapest to Supply

In practice, investors are often interested in targeting a given credit rating for the securities in their portfolio, say corresponding to some unconditional default probability, p^* . In light of the above discussion, which indicates that securities with the same default risk can trade at different prices, it is worth considering what would be the least expensive security that could be supplied to this investor, while satisfying his default risk target. Using the state-contingent pricing framework, one can show that within the set of discount bonds with an unconditional expected payoff of $1 - p^*$, the bond with the lowest price (i.e. the largest yield spread),⁵ will congregate the likelihood of default in the worst economic states possible. We refer to this security as the *cheapest to supply*.

Proposition 2 (CHEAPEST TO SUPPLY) *The cheapest security to supply with an unconditional expected payoff of $1 - p^*$ pays zero on a set with measure p^* containing the worst economic states, and one elsewhere.*

Proof: See Appendix A.

For example, if we were to order economic states, s , by the corresponding realization of the market return, the cheapest to supply bond would correspond to a digital call option on the market, with a strike price set at the p^* -th percentile of the τ -period market value distribution. At maturity, this simple call option delivers a payment of zero if the market value is in the bottom p^* percent of the anticipated realizations and one otherwise.

⁵Yield spread, or credit spread, refers to the annualized yield to maturity of a risky security in excess of the yield to maturity of a riskfree bond with similar duration.

As Figure 1 demonstrates, the payoffs of the digital tranche studied above increasingly resemble a digital call option as N becomes large. To see this, notice that in the limiting case of perfect diversification, the state-contingent portfolio payoff converges in probability to its mean, $1 - p_D(s)$. In turn, a tranche with a loss-attachment point of X only defaults if the realized economic state is below a critical threshold \tilde{s} , where $p_D(\tilde{s}) = X$. Thus, if we target default risk p^* , the payoff of the tranche referencing a fully diversified portfolio of economic assets converges to the payoff of the cheapest asset to supply with an unconditional default probability of p^* .

Proposition 3 *For any desired probability of default p^* , the cheapest asset to supply with that probability of default can be constructed by issuing a tranche with a loss attachment point of $X(p^*, \infty)$ against an asymptotically diversified collateral pool ($N \rightarrow \infty$). If the cumulative distribution of economic states is given by $\Pi(s)$, the attachment point of the limiting tranche is given by, $p_D(\Pi^{-1}(p^*))$.*

Proof: See Appendix A.

Finally, suppose the tranche was constructed to have the same unconditional default probability as the bonds in the underlying portfolio. In this case, the expected losses on the tranche and underlying portfolio are identical, and the two securities would command identical ratings. However, as N increases, the tranche converges to the cheapest asset to supply, causing it to trade at a yield spread that is greater than the yield spread on the underlying bond portfolio. This emerges vividly in Figure 2, which compares the payoffs of a bond portfolio and a digital tranche that have been constructed to have the same expected loss, as a function of the state of the economy. As we can see, relative to the underlying bond portfolio, the tranche has greater payoffs in good states (booms) and lower payoffs in bad states (recessions). Consequently, pooling and tranching emerges as a convenient method for altering the economic risk of the underlying securities, and, in particular, for manufacturing securities offering high yields relative to their default likelihoods. Investors who fail to appreciate the ability of pooling and tranching to reallocate payoffs across economic states of nature, may be attracted to the high yields offered by structured finance securities, without appreciating these securities' increased systematic risk.

2 Pricing CDO Tranches

One of the most widely issued structured instruments is the collateralized debt obligation or CDO. A CDO allows its originator to issue a prioritized capital structure of derivative claims against the underlying collateral pool. In a typical CDO, the underlying pool is either comprised of a portfolio of risky bonds (cash CDO) or credit default swaps (synthetic CDO), which can be thought of as default insurance contracts.⁶ This distinction, however, is immaterial for the derivation of the pricing model, as there is a simple no-arbitrage relationship linking the value of

⁶A credit default swap (CDS) gives the protection buyer the right to have the principal of an insured bond replaced at par in the event of default, in exchange for paying periodic insurance premiums to the protection seller.

a risky bond and its associated credit default swap. Consequently, in what follows, we focus on the valuation of cash CDOs, and subsequently show how to adapt our results to synthetic CDOs, which are the object of our empirical analysis.

The state-contingent valuation methodology described in the previous section can be easily applied to pricing real-world CDO tranches. To do so requires, (a) specifying a state space for pricing, and (b) modeling the expected payoffs for each tranche as a function of the realized economic state. In the spirit of the Sharpe (1964) and Lintner (1965) CAPM, we specialize to an economic state space defined by the realizations of equity index returns. This has the convenient feature that we are able to extract the corresponding state prices directly from market index options using the technique of Breeden and Litzenberger (1978).

To characterize the tranche payoffs in this economic state space, we need a model of the state-contingent payoffs of the underlying portfolio of risky bonds. With the state-contingent bond portfolio payoffs in hand, the CDO tranche payoffs can be simply obtained by noting that tranches are derivative claims, and their payoffs are determined contractually as a function of the payoff on the underlying portfolio. In order to characterize the state-contingent bond portfolio payoffs, we need a model to specify both the form of the state-contingent default probability, $p_D(s)$, and the state-contingent expected recovery rate, which will generally be different from zero. To do this, we modify Robert C. Merton's (1974) structural model of debt by imposing a factor structure on the asset returns of the firms issuing debt.⁷ Specifically, we assume asset returns satisfy a CAPM-style relation, which allows us to derive state-contingent expectations of the bond and tranche payoffs for all realizations of the market return. 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 correlation, which plays a key role in determining the distribution of losses on the underlying bond portfolio.⁸

Importantly, since we only rely on the model to produce payoffs conditional on the realization of the market return, we are able to remain agnostic about the distribution of the common factor. This represents a significant departure from existing single-factor models (Oldrich Vasicek (1987, 1991), Philipp Schönbucher (2000), Hull, Predescu and White (2006)), which make restrictive assumptions about the distribution of the common factor in order to derive closed-form expressions.⁹ Furthermore, because these models are estimated exclusively under the risk-neutral (i.e. pricing) measure, they obscure the important distinction between expected payoffs and risk premia, which can be transparently analyzed in our state-contingent valuation approach.

Finally, to value bond and CDO tranche payoffs we apply state prices extracted from equity index options. By using option-implied state prices, we achieve two important objectives. First, we ensure that we correctly capture the forward-looking distribution of the market return, which is

⁷See Young Ho Eom, Jean Helwege, and Jing-Zhi Huang (2004) and references therein, for a comprehensive survey of structural models.

⁸Chunsheng Zhou (2001) examines the ability of structural models to capture default correlations through asset correlations.

⁹The Vasicek model, with its normally distributed market factor, gives rise to a Gaussian copula for defaults, and is the basis of a common market quoting convention for CDO tranche prices, known as *implied correlation*.

likely to differ from the distributional assumptions commonly used in analytical models (Gaussian, student-t, etc).¹⁰ And, second, we capture the risk premia investors demand for assets which fail to pay off in adverse economic states, when the marginal utility of \$1 is high.

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

In Merton's model, a firm is assumed to default on its debt if the terminal value of its assets, $A_{i,T}$, falls below the face value of its debt, D_i .¹¹ Formally, since our goal is to model the expected payoffs for the portfolio of risky debt underlying the collateralized debt obligation (CDO), we would need to write down a separate model for each of the N firms, whose debt is in the reference portfolio. In order to avoid this complication, we instead make the simplifying assumption that the bonds in the portfolio were issued by N identical firms (*homogenous portfolio assumption*). Consequently, our model parameters are best thought of as characterizing the representative firm in the underlying portfolio. In turn, as in Merton's (1974) structural model, the representative firm will be characterized by a triple of parameters: the expected rate of return on its assets, which will be captured by the beta of its assets, β_a , the firm's debt-to-asset ratio, $\frac{D}{A_t}$, and the volatility of its assets, σ_a .

In order to obtain bond payoffs conditional on the realization of the market return, suitable for use in our state-contingent valuation method, we modify Merton's (1974) model by imposing a factor structure on the asset returns. Specifically, we assume that asset returns are driven by a combination of innovations to a market factor, \tilde{Z}_m , and idiosyncratic shocks, $\tilde{Z}_{i,\varepsilon}$. Importantly, since we only rely on the structural model to produce conditional payoffs, we leave the distribution of the market factor unspecified at this stage, allowing asset returns to have arbitrary, non-Gaussian unconditional distributions. Only *conditional* on the realization of the common factor, do we assume that asset returns are Gaussian (i.e. $\tilde{Z}_{i,\varepsilon}$ is normally distributed). Moreover, when writing down the returns for the firm's assets and the market, we focus on a simple, one-period model set in discrete-time, since defaults can only occur at maturity, T . If we note by τ , the time remaining to maturity, $T - t$, the (log) asset and market returns are assumed to be given by:

$$\ln \tilde{A}_{i,T} - \ln A_{i,t} = \left(r_f + \lambda_a - \frac{\sigma_a^2}{2} \right) \cdot \tau + \beta_a \sigma_m \cdot \sqrt{\tau} \cdot \tilde{Z}_m + \sigma_\varepsilon \cdot \sqrt{\tau} \cdot \tilde{Z}_{i,\varepsilon} \quad (5)$$

$$\ln \tilde{M}_T - \ln M_t = \left(r_f - \delta_m + \lambda_m - \frac{\sigma_m^2}{2} \right) \cdot \tau + \sigma_m \cdot \sqrt{\tau} \cdot \tilde{Z}_m \quad (6)$$

where: r_f denotes the riskless rate, λ_a is the asset risk premium, β_a is the asset CAPM beta and σ_ε is the idiosyncratic asset volatility, δ_m is the market dividend yield, λ_m is the equity market risk premium and σ_m is the market volatility. Finally, the total asset volatility, σ_a , is given by, $\sqrt{\beta_a^2 \sigma_m^2 + \sigma_\varepsilon^2}$, and the CAPM restriction on excess returns implies that: $\lambda_a - \frac{\sigma_a^2}{2} = \beta_a \cdot \left(\lambda_m - \frac{\sigma_m^2}{2} \right)$.

¹⁰Since our model relies on option prices to characterize the distribution of the common factor under the pricing measure, it can also be thought of as an option-implied copula model.

¹¹Fischer Black and John C. 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.

Since we are interested in deriving bond payoffs conditional on the realization of the common market factor, we will primarily be working with the conditional distribution of the terminal value of $\tilde{A}_{i,T}$. If we define the log moneyness, \tilde{m}_τ , as the logarithm of the ratio of the terminal market index level, \tilde{M}_T , to the time t futures price, $M_t \cdot \exp((r_f - \delta_m) \cdot \tau)$, one can show that the terminal asset value, $\tilde{A}_{i,T}$, conditional on terminal moneyness equaling m_τ , is a random variable given by:

$$\tilde{A}_{i,T}(m_\tau) = A_{i,t} \cdot \exp\left(r_f \cdot \tau + \beta_a \cdot m_\tau + \sigma_\varepsilon \cdot \sqrt{\tau} \cdot \tilde{Z}_{i,\varepsilon}\right) \quad (7)$$

Therefore the probability of a firm i defaulting on its debt conditional on the realization of m_τ , which is given by the probability that $\tilde{A}_{i,T}(m_\tau) < D$, is:

$$p_i^D(m_\tau) = \text{Prob}\left[\tilde{A}_{i,T}(m_\tau) < D\right] = \Phi[-\eta(m_\tau)] \quad (8)$$

where we have defined, $\eta(m_\tau)$ – the distance to default conditional on the realization of m_τ – as,

$$\eta(m_\tau) = -\frac{\ln \frac{D}{A_{i,t}} - (r_f \cdot \tau + \beta_a \cdot m_\tau)}{\sigma_\varepsilon \sqrt{\tau}}, \quad (9)$$

Unlike actuarial claims, whose default probability is unrelated to the economic state ($\beta = 0$), bonds are *economic assets* and have positive CAPM betas ($\beta > 0$). Consequently, their conditional probability of default increases in the adversity of the economic state ($\frac{dp_i^D(m_\tau)}{dm_\tau} < 0$).¹²

To derive an expression for the random, state-contingent bond portfolio payoff, we still have to specify a recovery value for any bonds that might have defaulted. In Merton's (1974) structural model, the payoff of a defaulted bond is naturally determined by the terminal value of the assets, $\tilde{A}_{i,T}$. However, it is now well-known that calibrated structural models generally imply counterfactually high unconditional recovery rates. Consequently, it is common to assume that a fraction, ν , of the terminal asset value is lost to bankruptcy costs (Hayne Leland (1994)). For example, in order to fit the data, Martijn Cremers, Joost Driessen, and Pascal Maenhout (2007) need bankruptcy costs to be approximately equal to 50% of the terminal asset value.¹³ With this auxiliary assumption, the random state-contingent payoff of an equally-weighted portfolio of N bonds – expressed as a fraction of the par value of the underlying bonds – is given by,

$$\tilde{P}_\tau(m_\tau) = \frac{1}{N} \sum_{i=1}^N \left(1 \cdot \left(1 - \tilde{\mathbf{1}}_{\tilde{A}_{i,T}(m_\tau) \leq D} \right) + \frac{(1 - \nu) \cdot \tilde{A}_{i,T}(m_\tau)}{D} \cdot \tilde{\mathbf{1}}_{\tilde{A}_{i,T}(m_\tau) \leq D} \right) \quad (10)$$

where $\tilde{\mathbf{1}}_{\tilde{A}_{i,T}(m_\tau) \leq D}$ is an indicator variable that – given the realization of \tilde{m}_τ – takes the value of one if bond i has defaulted, and zero otherwise.

¹²After conditioning on the realization of the market return, asset returns are multivariate Gaussian and uncorrelated, and thus independent. This implies that the distribution of the number of defaulted firms in the underlying portfolio of bonds will be binomial with parameter $p_D(m_\tau)$.

¹³As we will see in the next section, the average recovery rate (as a percentage of face value of debt) implied by this assumption is 40%, which is highly consistent with the empirical evidence of Edward Altman (2006).

To price the bond portfolio we need to compute its expected state-contingent payoffs and integrate them against the state prices extracted from equity index options. Since we have assumed that our portfolio consists of a homogeneous set of bonds, they will have identical state-contingent default probabilities and expected recovery rates. We therefore drop the i subscripts, and write the expected state-contingent bond portfolio payoff as

$$E_t \left[\tilde{P}_\tau(m_\tau) \right] = 1 - \left(1 - \frac{(1 - \nu) \cdot E_t \left[\tilde{A}_T(m_\tau) \mid \tilde{A}_T(m_\tau) < D \right]}{D} \right) \cdot p^D(m_\tau) \quad (11)$$

where $p^D(m_\tau)$ is the state-contingent default probability given in (8), and the expectation term reflects the conditional expectation of the terminal value of assets for a firm that has defaulted on its debt.¹⁴ The conditional expectation of a defaulted firm's terminal asset value is the mean of a truncated lognormal random variable and is equal to,

$$E_t \left[\tilde{A}_T(m_\tau) \mid \tilde{A}_T(m_\tau) < D_i \right] = A_t \cdot \exp \left(r_f \tau + \beta_a m_\tau + \frac{1}{2} \sqrt{\tau} \sigma_\varepsilon^2 \right) \cdot \frac{\Phi \left[-\eta(m_\tau) - \sigma_\varepsilon \sqrt{\tau} \right]}{\Phi \left[-\eta(m_\tau) \right]} \quad (12)$$

According to the above formula, the terminal asset value – which determines the ultimate bond recovery rate – is an increasing function of the market return, capturing the procyclical nature of recovery rates reported in the literature (Altman (2006)).

Finally, to price the underlying bond portfolio we apply the Arrow-Debreu valuation technique to its state-contingent payoffs. Since our state space is defined with respect to the realizations of the market factor, state prices can be readily computed from observations of equity index option values. In particular, Breeden and Lizenberger (1978) show that state prices can be extracted from European call option prices by taking the second derivative of the call price with respect to the strike at different moneyness levels. By integrating the product of the conditional expected portfolio payoff and the corresponding state price, $q(m_\tau)$, across all realizations of the market return, we arrive at the price (value) of the risky bond portfolio, V_τ :

$$V_\tau = \int_{-\infty}^{\infty} E_t \left[\tilde{P}_\tau(m_\tau) \right] \cdot q(m_\tau) dm_\tau \quad (13)$$

Since tranches are derivative claims on the underlying portfolio, their state-contingent payoffs are a function of the realized portfolio payoff in that state, $\tilde{P}_\tau(m_\tau)$. For example, consider a tranche with a lower loss attachment point of X and an upper loss attachment point of Y . The terminal payoff to this tranche is equal to 0 if the portfolio payoff is less than $1 - Y$; 1 if the portfolio payoff is more than $1 - X$; and is adjusted linearly between zero and one if the portfolio payoff is between $1 - Y$ and $1 - X$. This tranche payoff can be replicated exactly by combining a one dollar investment in a riskless bond, with a short position of $\frac{1}{Y-X}$ in put options on the bond portfolio

¹⁴Notice that if we set $\nu = 1$, such that recovery rates on defaulted debt are zero, the state-contingent expectation of the bond portfolio payoff becomes $1 - p^D(m_\tau)$, as in Section 1.

payoff struck at $1 - X$, and a long position of $\frac{1}{Y-X}$ put options struck at $1 - Y$:

$$\tilde{P}_\tau^{[X,Y]}(m_\tau) = 1 - \frac{1}{Y-X} \cdot \left(\max\left((1-X) - \tilde{P}_\tau(m_\tau), 0\right) - \max\left((1-Y) - \tilde{P}_\tau(m_\tau), 0\right) \right) \quad (14)$$

Conceptually, in order to obtain the corresponding prices for the tranches, $V_\tau^{[X,Y]}$, one proceeds in the same way as in the case of the underlying portfolio, by first computing the expected state-contingent payoffs and then integrating them against the state prices, as in (13). However, in the case of the tranches the expected state-contingent payoffs are difficult to compute analytically because of the complexity of $\tilde{P}_\tau(m_\tau)$ and the non-linearity of the tranche payoff function, (14). Consequently, it is necessary to resort to numerical methods in order to evaluate the payoff expectation appearing in the pricing integral. In practice, the expected state-contingent tranche payoff can be obtained with little computational effort, by simulating the conditional payoffs of the underlying bond portfolio and applying the tranche contract terms.

3 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 Francis Longstaff and Arvind Rajan (2007), consists of an equally-weighted portfolio of 125 liquid five-year credit default swap (CDS) contracts on U.S. firms with investment grade corporate debt.¹⁵ A CDS is essentially a default insurance contract, which allows two parties to transfer the credit risk of the reference entity. The contract obligates the protection buyer to pay a periodic fee, known as a running spread, on the notional amount being insured to the protection seller, in exchange for having the protection seller assume the contingent liability of replacing the par value of the bond in case of default. The running spread on the CDX can be thought of as the cost of insuring a pre-specified notional amount of an equally-weighted portfolio of risky bonds. It can be represented as a weighted average of the spreads on the individual CDSs, with weights that adjust for the relative riskiness of the payment streams to be received on each of the contracts. When the underlying securities are reasonably homogenous in terms of their default risk, the CDX spread will be approximately equal to the average of the underlying CDS running spreads.

Our data cover the period September 2004 to September 2007, and contain daily spreads of the CDX as well as spreads on the 0-3, 3-7, 7-10, 10-15, and 15-30 CDX tranches.¹⁶ The CDX tranches

¹⁵Darrell Duffie and Kenneth J. Singleton (2003) provide a textbook treatment of credit default swaps (CDSs) and collateralized debt obligations (CDOs). Additional surveys of credit derivatives are provided by Darrell Duffie and Nicolae Gârleanu (2001) and references therein.

¹⁶An earlier version of this paper, first circulated in June 2007, was based on credit market data through September 2006. Beginning in August 2007, credit markets entered a period of turmoil during which market conditions departed sharply from the earlier sample. The opportunity to confront our model with data that was from a different market regime, prompted us to extend the time series through September 2007 to produce the dataset that is the focus of the current version. The earlier working paper is available from <http://hbswk.hbs.edu/item/5725.html>.

are derivative securities with payoffs based on the losses on the underlying CDX, and are defined in terms of their loss attachment points. For example, a \$1 investment in the 7-10 tranche receives a payoff of \$1 if the total losses on the CDX are less than 7%; \$0 if total CDX losses exceed 10%; and a payoff that is linearly adjusted for CDX losses between 7% and 10%. As with CDSs, the tranche prices are quoted in terms of the running spreads that a buyer of protection would have to pay in order to insure the tranche payoff.¹⁷ In the case of tranches, the protection buyer pays the running spread only on the surviving tranche notional on each date. Consequently, if the portfolio losses exceed the tranche’s upper attachment point, the protection buyer ceases to make payments to the protection seller.

In practice, the composition of the CDX index is refreshed every March and September to reflect changes in the composition of the liquid investment grade bond universe. In turn, each new version of the CDX, referenced by a series number, remains “on-the-run” for six months after the roll date. Since the majority of market activity is concentrated in the on-the-run series, we splice the first six months of Series 3 through Series 8 of the CDX NA IG to produce a continuous series of on-the-run spreads over the three-year period, as in Longstaff and Rajan (2007). For contextual reference, we compare the CDX and tranche credit spreads to daily series of average corporate bond spreads on AA, A, BBB, BB, and B-rated bonds. These spreads are reported in terms of an equivalent 5-year CDS spread implied by corporate bond prices.

Finally, our analysis also requires accurate prices for out-of-the-money market put options in order to construct state prices for use in valuation. To match the timing of the CDX payoff, we require put options that have five years to maturity. Unfortunately, during our sample period, no index options with maturity exceeding three years traded on centralized exchanges. Instead we rely on daily over-the-counter quotes on five-year S&P 500 options obtained from Citigroup. These quotes correspond to 13 securities with standardized moneyness levels ranging from 0.70 (30% out-of-the-money) to 1.30 (30% in-the-money) at increments of 5%, allowing us to fit an implied volatility function for long-dated options on each day.¹⁸

3.1 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 standard deviations for each of our series across the sample. As expected, average spreads are increasing across the bond portfolios and across the tranches as the credit quality falls. The average CDX spread is very close to that of the BBB index, consistent with the average rating of the individual CDS contracts in the CDX being rated BBB, as reported by Atish Kakodar and Barnaby Martin (2004).

¹⁷Formally, the market quoting convention for the 0-3 tranche is to quote an up-front payment to be received by the protection seller at the initiation of the contract, in addition to a fixed running spread of 500bps. In order to make the 0-3 tranche price series comparable to the remaining tranches, we convert the up-front payment into an equivalent running spread.

¹⁸An option’s *moneyness* is defined as $\frac{K}{F_{t,\tau}}$, where K is the option strike price and $F_{t,\tau}$ is the contemporaneous τ -year ($\tau = 5$) futures price on the S&P 500 index.

Panel 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 changes in all bond spreads, suggesting that market volatility is a key factor in the pricing of credit-sensitive corporate securities.¹⁹ The correlation between changes in long-term volatility and the changes in the CDX and the tranche spreads is somewhat larger.

4 Calibrating the Bond Pricing Model

A seller of protection on the CDX is effectively committing to bearing the losses on the underlying portfolio of risky debt in exchange for receiving periodic spread payments. In order for no payments to exchange hands at contract initiation, the quoted CDX spread is set such that the present value of the spread payments to be made over the life of the contract by the protection buyer (premium leg) is exactly equal to the present value of the expected losses to be borne by the protection seller (protection leg). Consequently, while our ultimate goal is to match the observed CDX spread, the above relationship allows us to focus on modeling the present value of expected losses, which we then convert into an equivalent running spread. The present value of the losses, however, is simply given by the wedge between the value of a riskless debt portfolio and the risky debt portfolio underlying the CDX, V_τ , which can be priced using the model in Section 2.

In order to price the underlying risky debt portfolio referenced by the credit default swaps in the CDX, we adopt a simple state-contingent valuation approach. In particular, we use the structural model in Section 2 to produce the expected state-contingent payoffs for the underlying portfolio, (11), which we then value using an empirical estimate of the state prices obtained from 5-year index options, as in (13). In other words, we project the payoffs on the risky debt into the space of market returns using the structural model, and then use Arrow-Debreu prices to arrive at the price of the risky debt portfolio. The price is then converted into an equivalent spread that would be paid by an agent seeking to insure the portfolio against potential losses.

To simplify the calibration procedure we make the auxiliary assumption that the risky bond portfolio underlying the CDX is homogenous, i.e. it is comprised of bonds issued by N identical firms. Consequently, rather than estimate a triple of parameters – the firm leverage ratio, $\frac{D}{A_t}$, the firm’s idiosyncratic asset volatility, σ_ε , and the firm’s asset beta, β_a – for each firm in the pool, we only estimate one set of parameter values, which are best thought of as characterizing a representative firm in the index. On each day, to pin down the three parameters, we require that the following three constraints be satisfied: (1) the model-implied CDX running spread matches the empirically observed running spread; (2) the model-implied equity beta of the firm is equal to $\bar{\beta}$; and, (3) the model-implied pairwise equity return correlation is equal to $\bar{\rho}$.

An implicit assumption of this calibration procedure, consistent with industry practice, is that the CDX spread reflects the risk-adjusted compensation for the expected loss given default, and

¹⁹Jun Pan and Kenneth J. Singleton (2006) find that sovereign credit default swap spreads move with measures of aggregate volatility.

is unaffected by tax or liquidity considerations. Indeed, 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 than the corporate bond market.

The two ingredients of our model calibration are the state-contingent payoff function – parameterized by $(\frac{D}{A_t}, \sigma_\varepsilon, \beta_a)$ – and the option-implied state prices. If the structural model correctly captures the risk characteristics of the underlying portfolio, the calibrated set of parameter values should allow us to not only match the level of the CDX on a given day with sensible parameter values, but also help explain variation in the CDX spread as state prices change. To verify the robustness of our model along these lines, we perform a variety of checks. First, we compare the performance of two parametric implied volatility functions used in constructing the state price density. Second, we evaluate the performance of two recovery rate assumptions. Third, we show that our calibration procedure allows us to obtain high R^2 in forecasting CDX yield changes at various frequencies, for various combinations of implied volatility specifications and recovery assumptions. 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 calibrated model can be used to price tranches, as well as, inform the construction of simple replicating strategies involving put spreads on the market index.

4.1 Extracting State Prices

In order to value the state-contingent bond portfolio payoffs produced using model in Section 2, we need a complete set of state prices as a function of the realized τ -period market return, $q(m_\tau)$. As before, we identify states by their (log) moneyness, which is defined as the log ratio of terminal market value, $M_{t+\tau}$, to the τ -period futures price at time t ,

$$m_\tau = \ln \frac{M_{t+\tau}}{F_{t,\tau}} = \ln \frac{M_{t+\tau}}{M_t \cdot \exp((r_f - \delta_m) \cdot \tau)} \quad (15)$$

To extract state prices on this grid we exploit the fact that the prices of Arrow-Debreu securities can be recovered from options data. In particular, Breeden and Litzenberger (1978) have shown that – given the market prices of European call options with τ -periods to maturity and strike prices K , $C(K, \tau)$ – the price of an Arrow-Debreu security for state m_τ is equal to the second derivative of the call price function with respect to the strike price. Consequently, the state price for state m_τ can be computed from,

$$q(m_\tau) = \frac{\partial^2 C(K, \tau)}{\partial K^2} \Big|_{K=x \cdot F_{t,\tau}} \quad (16)$$

where $x = \exp(m_\tau)$ denotes the corresponding moneyness level. The formula for the Arrow-Debreu prices is particularly simple when the underlying asset 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, if we express option prices, $C(K, \tau)$, using the Black-Scholes formula and allow implied volatility to be a function of the strike price, $C^{BS}(K, \sigma(K, \tau), \tau)$, we obtain,

$$q(m_\tau) = \frac{\partial^2 C^{BS}}{\partial K^2} + \frac{d\sigma}{dK} \cdot \left(2 \cdot \frac{\partial^2 C^{BS}}{\partial K \partial \sigma} + \frac{\partial^2 C^{BS}}{\partial \sigma^2} \cdot \frac{d\sigma}{dK} \right) + \frac{d^2 \sigma}{dK^2} \cdot \frac{\partial C^{BS}}{\partial \sigma} \Big|_{K=x \cdot F_{t,\tau}} \quad (17)$$

For any twice-differentiable implied volatility function, $\sigma(K, \tau)$, the price of the corresponding Arrow-Debreu securities can be computed in closed-form using the above expression. Consistent with intuition, the state prices 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. When the implied volatility function is flat – as would be the case when the underlying asset follows a log-normal diffusion – only the first term remains.

Since any twice-differentiable implied volatility function implies a set of Arrow-Debreu prices, the challenge of computing state prices effectively boils down to fitting an implied volatility function to the empirical option data. Of course, since we only have observations for a discrete set of options with moneyness values from 0.7 to 1.3, we are forced to both interpolate and extrapolate the implied volatility function. Typical parametric approaches to this problem involve fitting linear combinations of orthogonal polynomials (Joshua Rosenberg and Robert F. Engle (2001)) or splines (Robert Bliss and Nikolaos Panigirtzoglou (2004)).²⁰ Instead, we take a more direct route, and propose a few parsimoniously parameterized implied volatility functions, which produce strictly positive implied volatilities, have controllable behavior in the tails, and are twice differentiable (see Appendix B for details). Our preferred specification is based on the hyperbolic tangent function and takes the form:

$$\sigma(x, \tau) = a + b \cdot \tanh(-c \cdot \ln x) \quad (a > b > 0) \quad (18)$$

where x denotes the moneyness level. To select the optimal parameter values for the proposed implied volatility specification, we compute the corresponding Arrow-Debreu prices, value the thirteen European options for which we observe prices, and minimize the sum of squared (percentage) pricing errors. This procedure pins down the optimal parameter values for the proposed specification, allowing us to extrapolate the implied volatility on the entire grid of m_τ , and compute a complete set of Arrow-Debreu state prices from (17) for use in valuation. The resulting daily state price densities are highly consistent with the observed option values, producing an average daily root mean squared percentage pricing error across all securities of 0.29%.

Figure 3 displays the calibrated 5-year state prices and implied volatility functions for three CDX initiation dates. The average 5-year at-the-money implied volatility is around 20% and is

²⁰Yacine Ait-Sahalia and Andrew W. 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 Jens C. Jackwerth (1999) or Bernhard Brunner and Reinhold Hafner (2003).

about 10% at very high moneyness levels. At a moneyness of 0, the implied volatility averages roughly 30%. The implied state price densities tend to have very fat left tails between moneyness levels of 0 to 0.5, reflecting the high price of bad economic states expressed in the index options market.

The ability to construct a set of Arrow-Debreu prices spanning all possible realizations of m_τ , even those for which we do not have empirical option data, is crucial to our analysis. The key concern is that our parametric implied volatility function may over-estimate the state-prices associated with bad economic states, which would then overstate our estimates of the required risk compensation for the CDX tranches. To mitigate this possibility, we also constrain the implied volatility function to have a maximum implied volatility equal to the maximum implied volatility observed among the option prices for which we have data. This essentially amounts to flattening the implied volatility skew between moneyness of 0 and 0.7 at the level observed at moneyness of 0.7. This alters the state price density so that the pricing errors associated with the observed options are somewhat larger,²¹ but with the benefit of producing conservative deep out-of-the-money state prices that can be used to evaluate the robustness of our tranche prices. We refer to model specifications using this constrained state price density as having a “constrained implied volatility function.”

4.2 Implying the Conditional Payoff

Equation (13) shows that the value of the risky debt portfolio referenced by the CDX can be simply obtained by scaling the portfolio’s state-contingent payoffs by the state prices, and summing across economic states, m_τ . 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, Driessen, and Maenhout (2007), which finds that the pricing of individual credit default swaps is consistent with the option-implied pricing kernel.

Given our empirical estimates of the daily set of state prices, our actual observations of the CDX price, and a structural model of the conditional payoff, we can calibrate the model parameters $\left(\frac{D}{A_t}, \beta_a, \sigma_\varepsilon\right)$ to match the CDX spread on each day. To find a unique solution, we impose the two additional constraints on the average firm beta and correlation described above. In particular, we require that the model implied equity beta and pairwise correlation match their empirical counterparts. We estimate the average pairwise equity correlation, $\bar{\rho}$, for the constituents of the CDX over the period 2003 to 2007 to be 0.2 and the average equity beta, $\bar{\beta}$, to be one. This is not surprising since the CDX is comprised of investment grade securities issued by some of the largest U.S. corporations.

Our baseline payoff specification for the risky debt portfolio, (11), relies on the Merton model recovery rate, which has a strong procyclicality. One potential concern is that this procyclicality in recovery rates implies too much impairment in bad economic states, which will overstate the

²¹The constrained implied volatility function results in an average root mean squared percentage error across all securities of 3.20%.

required risk compensation. To mitigate this concern, we also evaluate the model with an alternative recovery rate assumption. Specifically, we calculate expected payoffs using the Merton model default probability and a simple state-independent mean recovery rate, \bar{R} , of 40%:²²

$$E_t \left[\tilde{P}_\tau(m_\tau) \right] = 1 - (1 - \bar{R}) \cdot p^D(m_\tau). \quad (19)$$

With this assumption, the conditional payoff of the CDX goes to 0.40 as the market index goes to zero, whereas with the Merton model recovery assumption the conditional CDX payoff goes to zero as the market index loses all of its value. An advantage of this approach is that, because default is independent of firm and market value, it offers a conservative, downward-biased estimate of the amount of systematic risk in the underlying bonds.

We then perform daily calibrations of the firm parameters, $\left(\frac{D}{A_t}, \beta_a, \sigma_\varepsilon \right)$, under the four models that obtain from our two parametric implied volatility specifications and our two recovery rate assumptions. This results in a time series of the underlying firm parameters and implied measures of default probability and recovery rates under the pricing measure. Summary statistics of these parameters and implied values are reported in Table 2. Across all four model specifications, the mean parameter values are highly similar, as are the implied credit risk measures. This suggests that the specific parametric assumptions about the state price density and the expected payoff function in the deep out-of-the-money states are not pivotal to explaining the CDX spread. Moreover, the parameters exhibit little time series variation, suggesting that the parameters themselves are quite stable through time.

4.3 Evaluating the Bond Pricing Model

Since the calibration procedure identifies a triple of firm parameters using three constraints, it enables the model to match the level of the CDX spread on each day in our sample. Therefore, in order to assess the model’s ability to accurately characterize the priced risks of corporate bonds, we turn to the time series dynamics of the spread on the model-implied replicating portfolio of Arrow-Debreu securities. If the model correctly matches the risk characteristics of the underlying securities, the change in the yield on the model-implied replicating portfolio should match the actual dynamics of the CDX spread. Consequently, 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 3 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 model yield at time t . The model predicted change is highly statistically significant with a large R^2 for all implementations of the model. In the baseline specification (Model 1), the

²²The CDS-implied spread data we are using assumes a 40% recovery value which is also consistent with Altman (2006). We obtain qualitatively similar results if we use a 50% recovery rate.

model predicted change has a t -statistic of 8.70 and an R^2 of 0.34.

We also examine the relation between spread changes and changes in the riskfree rate, the S&P 500 index level, the 5-year at-the-money implied volatility, and the 5-year implied volatility skew. These variables are each statistically significant in explaining changes in the CDX spread, but largely lose significance when the model predicted change is included. The exception is the change in the implied volatility skew, which continues to remain statistically significant. 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. Again, differences in the treatment of the left tails of the state price density and the expected payoff function appear to have little affect on the model's ability to explain changes in CDX spreads as evidenced by the qualitatively similar results across all specifications. Finally, the explanatory power of the model compares favorably to other empirical investigations into the determinants of credit spread changes for corporate bonds and CDSs (Pierre Collin-Dufresne, Robert S. Goldstein, and J. Spencer Martin (2001) and Benjamin Zhang, Hao Zhou, and Haibin Zhu(2006)).

5 Pricing Credit Derivatives

The Merton (1974) credit model integrated with a common market factor, in the spirit of 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 prices. In other words, this pricing framework provides a direct link between the bond market and the equity 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. We now turn to the question of whether the prices of tranches issued on the bond portfolio are consistent with their market risks.

This unified framework makes pricing credit 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 using (10).²³ Since the tranche (i.e. derivative) payoffs are defined as a function of the payoff on the CDX (i.e. the underlying security), they can be identified by simply applying the contract terms to each simulated outcome. In this case, each tranche's payoff is defined by its loss attachment points to the CDX payoff. Finally, pricing is completed as before, by computing mean state-contingent payoffs and applying the Arrow-Debreu prices.

The simulation of the state-contingent tranche payoffs proceeds as follows. First, we specify a moneyness grid (i.e. possible market realizations in terms of gross return), ranging from 0 to 10 at increments of 0.005. Second, for each day in the sample, at each point along the grid, we make use

²³Model specifications in which the mean recovery value is state-independent require an additional assumption about the conditional distribution of the firm-level loss. In particular, 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 25%, based on empirical estimates of these values from Altman (2006).

of equation (10) and the calibrated firm parameters to simulate the terminal values for each of the firms in the CDX ($N = 125$), which determine the portfolio loss. The terms of each tranche are applied to the portfolio loss to calculate the various tranche payoffs. We repeat this step 10,000 times and then calculate the state-contingent mean payoff for each tranche, which is valued using that day’s option-implied state prices.²⁴ The simulated state-contingent payoffs for the CDX and the 7-10 tranche, and their associated state-contingent means, are displayed in Figure 4.

Table 4 presents a comparison of the spreads predicted by the models with the spreads offered by each of the CDX tranches. In particular, for all implementations of the model, we report the time series mean of the actual and model spreads and the correlation between weekly yields and changes in yields. The first-loss, or equity, tranche (0-3) has a higher mean spread than the model predicts, suggesting that this security has been undervalued in the market relative to our model.²⁵ Our model predicts considerably greater spreads than are present in the data for all of the other non-equity tranches. As a fraction of the spread, the disparity is most severe for the relatively senior 7-10 and 10-15 tranches. The 7-10 tranche spread predicted by our model exceeds actual spreads by more than a factor of three. For the 10-15 tranche our model predicts spreads that are four times as large as in the data. Figure 5 graphs the predicted and actual spreads through time. As can be seen, the model spreads exceed actual spreads across the entire sample for each of the non-equity tranches. 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 credit spread levels are off by a large amount, the returns offered by the model tranches and their actual counterparts are driven by common economic risks.

The formal Arrow-Debreu pricing approach assumes that markets are complete. In our setting, this translates into assuming that there is a continuum of options available with strike prices for all possible realizations of the market index. The replicating portfolio implicit in the formal model essentially takes positions in all of these securities. A simpler and highly feasible replicating strategy is to approximate the mean state-contingent tranche payoff with a portfolio comprised of a riskless bond and a buy-and-hold position in a single put spread on the market index (i.e. short an index put option with a relatively high strike price and long an index put option with a relative low strike price) with appropriately chosen strike prices. This replicating portfolio has two primary benefits. First, it relies only on two options rather than the continuum of options implicit in the model. Second, because this replicating strategy represents a static portfolio with no rebalancing for 6 months, it is attractive from an implementation perspective.

The challenge is to determine what index option strike prices should be used to replicate the market exposure for each tranche. While there are several ways to develop this approximation,

²⁴A convenient approximation to the simulation procedure is to simply evaluate the tranche payoffs at the *expected* state-contingent CDX payoff, rather than computing their expectation as a function of the random, state-contingent CDX payoff. Although this approach neglects the Jensen correction terms stemming from the non-linearity of the tranche payoff function, it is significantly faster and results in empirically negligible errors.

²⁵The conversion of the upfront spread to a running spread for the 0-3 tranche requires estimating the timing and magnitude of expected CDX losses, which is model specific. This causes the “actual” 0-3 tranche spread reported in Table 4 to vary slightly across model specifications.

we simply set the strike prices by matching the tranche attachment points to the *expected* CDX payoff. For example, in order to find the strike prices of the puts corresponding to a tranche with a lower attachment point of X , and an upper attachment point of Y , we set (11) equal to $1 - X$, and solve for the log moneyness level, m_τ . Conveniently, the exponential of the resulting log moneyness solution yields the strike price of the less out-of-the-money index put option (i.e. the option to short). To find the strike price of the more out-of-the money index put option (i.e. the option to go long), we repeat the computation setting the expected CDX payoff equal to $1 - Y$.

Table 5 compares the actual spreads for each tranche to the market-risk-matched put spreads. The results using the static put-spread approximations are very similar to those of the formal model.²⁶ Again, the replicating portfolios offer considerably larger yield spreads than those available from the non-equity CDO tranches. Moreover, the correlations between weekly changes of the actual credit spreads and the put spread approximate model are quite high. Taken together, it appears that senior CDX tranches have offered too little compensation for the market risks to which they are exposed when compared to the compensation investors were able to earn in other markets for bearing similar market exposures.

6 Discussion

Our paper begins with the theoretical prediction that senior CDO tranches concentrate expected losses in adverse economic states relative to single-name corporate bonds, and the observation that this theoretical prediction has not been well-articulated in either the academic or practitioner literatures. This creates the possibility that some market participants may not fully appreciate that the risk exposures of structured finance products are fundamentally distinct from single-name corporate bonds, even though their credit ratings can be directly compared. If these participants have an impact on prices, the CDO market will appear inefficient, with senior tranches offering insufficient compensation relative to their underlying risks.

Tests of market efficiency are inherently joint tests of efficiency and a pricing model (Fama (1970)). Therefore, in interpreting our empirical finding that senior CDX tranches indeed appear to have been overpriced, it is useful to review the research design and explore the robustness of the results to alternative model specifications.

6.1 Model Robustness

The empirical analysis in this paper rests on three key assumptions: (1) a single market factor explains all common movements in asset returns; (2) our sample of index options is sufficient to interpolate and extrapolate to all relevant state prices; and (3) the constituents of the CDX are sufficiently numerous and homogeneous to be accurately modeled by a representative firm. This section will explore and discuss the robustness of our results to each of these assumptions.

²⁶It turns out that the discrepancy between the formal model and the static market index put spread approximation is almost entirely due to the difference in rebalancing frequencies.

6.1.1 Single-Factor

A 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 bode more poorly for bondholders than an equivalent, firm-specific decline.

Figure 6 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. Overall, 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.

6.1.2 Implied Volatility Extrapolation

The second key assumption of our paper is that our sample of five-year over-the-counter implied volatilities is sufficient to interpolate and extrapolate to all relevant state prices. In particular, our volatilities only cover the moneyness range of 0.7 to 1.3. Many of our tranches are highly exposed to states associated with moneyness levels that are well below 0.7. Therefore, to provide additional perspective on the likely shape of the 5-year implied volatility function for moneyness levels below 0.7 and to verify that our results are not an artifact of the OTC volatility prices we obtained, we examined long-dated exchange-traded options on the S&P 500 index (SPX). Specifically, we identified on each date the longest maturity SPX options and calculated their implied volatilities across all observed strike prices. These options are on average two and a half years to maturity and almost always include options with moneyness levels that are below 0.7, often as low as 0.4. The exchange-traded options have volatility levels and skews that are highly similar to the five-year volatilities used in our analysis. More importantly, the exchange-traded options continue to exhibit a monotonic volatility skew for all moneyness levels below 1.0. Indeed, during our sample period,

the skew below moneyness of 1.0 in exchange-traded implied volatilities was strictly downward sloping on every day.

We also examined a second, far more conservative implied volatility specification which restricted the implied volatilities below a moneyness of 0.7 to lie strictly below the implied volatility at a moneyness of 0.7. As discussed in Section 4, this flattened implied volatility function does not impact our results in a material way. And finally, in unreported results, we used an exponential implied volatility function of the form $\sigma(x, \tau) = a + b \cdot \exp(-c \cdot x)$, where x is the option moneyness. Because this form is uniformly convex and therefore assigns extremely high implied volatilities for deeply out-of-the-money options, it leads to an even stronger mispricing. In fact, during the summer of 2007 the observed volatility skew becomes sufficiently steep that estimates of the out-of-the-money volatility are too large to be reconciled with the CDX index level, let alone the tranches.

6.1.3 Homogeneity

A final assumption relied upon in our calibration exercise is that the dynamics of the 125 firms in the CDX could be accurately captured by a single representative firm. First, it is worth noting that our approach explains CDX movements with R^2 values that are at least as high as those elsewhere in the literature, suggesting that this assumption is not completely at odds with alternative approaches.

Second, and somewhat related to the single-factor assumption described above, a potentially important aspect of the cross section of underlying securities is within-industry return correlation. To the extent that this represents a significant source of co-movement, it could change the relative pricing of the equity and mezzanine tranches, as greater co-movement increases the likelihood of a mezzanine default. On the other hand, it is important to remember that industry co-movement will be, by design, orthogonal to market movements. And, because the CDX is well-diversified pool across industries, even the default of an entire industry is unlikely to create portfolio losses in excess of the seven percent that is needed to impair the senior tranches. As a result, it is difficult to see how the introduction of within-industry correlation can alter the risks faced by tranches with senior levels of subordination in a material way. Relatedly, some evidence on the importance of industry-level correlations comes from the May 2005 downgrade of GM and Ford, which clearly reflected industry-level correlation in their respective cashflows. This event ultimately led to an increase in mezzanine tranche prices and a downward reassessment of the importance of industry shocks. Equity tranche investors who hedged their positions by shorting the mezzanine tranches experienced significant losses during this period as their models had overestimated the ability of an industry-level shock to wipe out the equity tranche and impair the mezzanine tranche.

6.2 Ruling Out Extreme States

The basic insight behind this paper is that in adverse economic states, losses on CDX tranches are expected to exceed those of single-name corporate bonds with an otherwise identical uncon-

ditional expected loss, or credit rating. Of course, these are also economic states that we do not empirically observe very often, if ever. This raises the concern that our tranche pricing results are driven by the interaction of high, but imprecisely estimated, state prices in unobserved economic states (i.e. deep out-of-the-money states) with large expected losses in these same states.

Since most models will struggle to produce sufficiently precise estimates in regions never before observed, this issue can only be addressed by stepping outside of the model. Rather than estimate state prices for deeply out-of-the-money states for which we lack data, we make the extreme counterfactual assumption that state prices in this region are equal to zero. This explicitly introduces an arbitrage in the index option market, in that insurance against positive probability events can be purchased for free, but strongly biases us against finding a mispricing in CDX tranches. In particular, we assume that the state prices below moneyness of 0.40 are zero, and shift the mass associated with these states proportionally to all other states.²⁷ This alteration of the state-price density rules out the possibility of observing a terminal market value below 0.40 of the prevailing futures price, and results in a dramatic rise in the option pricing errors for the observed index options, especially for out-of-the-money put options. Notably, this proportional redistribution of the state price density is conservative relative to one that reallocates this mass with the objective of minimizing pricing errors of the observed options. We then repeat the daily CDX calibration exercise and examine the various CDX tranche spreads (see Figure 7). Interestingly, after making this extraordinary adjustment to the state prices, our model continues to suggest that the equity tranche has been underpriced, while the 3-7, 7-10, and 10-15 tranches have been overpriced. This highlights that the main results for all the tranches other than the super senior tranches are not driven by the most extreme economic states.

6.3 Which Market is Mispriced?

Our empirical analysis represents a relative pricing exercise in which tradable prices are compared across three distinct markets: the S&P index options market, the corporate bond market, and the CDX tranche market. Although we have interpreted our results as suggesting that senior CDO tranches are grossly overpriced, a natural question is whether we can rule out the possibility that the CDX tranches are correctly priced and one of the other two markets is mispriced. This question is of particular interest in light of the literature questioning the efficiency of the index options market.²⁸ In addition to circumstantial evidence that suggests that the mispricing can be attributed to the CDO market, including the extremely rapid growth in CDO issuance over the past several years and the fairly substantial repricing of senior CDO tranches that began at the end of our sample and has continued well into 2008, our structural approach points in this direction as

²⁷This adjustment is made to the baseline state price density based on the hyperbolic tangent implied volatility function. The 0.40 moneyness cutoff is chosen on the basis of the lowest observed moneyness among exchange-traded equity index options.

²⁸Oleg Bondarenko (2003) and David Bates (2006), for example, reach the conclusion that index options are overpriced, whereas Mark Broadie, Mikhail Chernov, and Michael Johannes (2008) argue that observed index options data are not inconsistent with standard models.

well.

A feature of our model is its ability to reconcile the prices of index options and the underlying bond portfolio. The calibration delivers stable and sensible parameter estimates and does a respectable job in explaining bond price dynamics through time. We interpret this as evidence that the bond and option markets are reasonably well-integrated. An alternative exercise is to ask whether our model can reconcile the bond and tranche prices. For example, recent papers by Longstaff and Rajan (2007), Shahriar Azizpour and Kay Giesecke (2008), and Andreas Eckner (2008), find that CDX and CDX tranche prices can be made consistent with reduced-form models and conclude that these securities are reasonably efficiently priced.

In the context of our setup, this amounts to asking whether there exists a reasonable set of option prices, equivalently a state price density, that jointly fits the spreads on the CDX and CDX tranches. To investigate this possibility, we search over the set of option prices that can be characterized by the hyperbolic tangent implied volatility function. This approach essentially allows us to vary the level and slope of option implied volatilities as a function of moneyness. Within this class of implied volatility functions, we are unable to find a set of option prices that can jointly price the CDX and CDX tranches. State price densities required to match the CDX underprice the senior tranches and those that match the tranche prices overprice the CDX. Based on this, we conclude that our model only permits two interpretations of the data: either CDO tranches are mispriced or *both* index options and corporate bonds are mispriced.

6.4 Credit Risk Premia

As we have shown in Section 1, theory predicts that securities with identical credit ratings – assessed on the basis of expected probabilities of default, or expected losses – can trade at different prices. In particular, creating large diversified portfolios of economic assets (e.g. corporate bonds) and issuing prioritized capital structures of claims against those pools, as is common in structured finance, emerges as a natural approach to manufacturing the cheapest security within a given credit rating category. When the underlying portfolios become asymptotically diversified ($N \rightarrow \infty$), the resulting claims concentrate the likelihood of default in the most adverse economic states, effectively creating economic catastrophe bonds. In order to compensate investors for the systematic risk they are bearing, correctly priced ECBs offer the maximum possible yield spread per unit of default risk. To capture this bound, and more generally, the variation in compensation per unit of default risk, we introduce the credit risk ratio, $\phi = \frac{Yield\ Spread}{Loss\ Rate}$.²⁹ While the yield spread reflects both compensation for expected losses and risk premia arising from covariation of losses with economic outcomes, the loss rate is an annualized spread over the riskless rate necessary to compensate the investor for the risk of loss with no risk premium.³⁰

²⁹When recovery rates are zero, ϕ collapses to the ratio of the risk-neutral and actual default intensities – a commonly-used statistic in the reduced-form credit literature.

³⁰The loss rate is given by,

$$Loss\ Rate = -\frac{1}{T} \ln E[\text{Payoff}]$$

While inconsistent with the prices of market index options, the assumption of a lognormal distribution for the terminal market value is helpful for further assessing some of the model’s implications and comparing estimated loss rates of the tranches to the rating-based bond indices. In particular, with an explicit distribution of market states we can decompose credit spreads into compensation for expected loss and systematic risk. Therefore, to compute loss rates we make an auxiliary assumption that the terminal distribution of the market is lognormal, with a market risk premium, λ_m , equal to 5% per year and volatility, σ_m , given by the at-the-money 5-year option-implied volatility. This also allows us to calculate unconditional default probabilities and average values of ϕ for each of the securities. These values are reported in Table 6.

The credit risk ratio reflects the relative importance of the risk premium in the pricing of a defaultable bond, and is related to the average value of the marginal utility in states in which the bond is likely to default. For example, if a bond’s defaults are idiosyncratic, 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 experience losses in states with high marginal utility the higher the value of the ratio. Consistent with intuition and previous empirical findings, credit risk ratios are increasing in the seniority of the rating-based bond indices.³¹ This indicates that the average economic state in which a highly rated bond defaults is worse than the average economic state in which a lower 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. The mean calibrated default intensity for the CDX is 31 bps, corresponding to a 5-year default probability of 1.54%, which is between that for A-rated (0.50%) and BBB-rated (2.08%) bonds, as reported in Cantor et al. (2005). The credit risk ratio for the CDX averages 2.9, and is fairly stable through time, ranging from 2.3 to 3.3.

Recent structured finance activity has created a new generation of products by pooling senior corporate claims (investment grade bonds in this paper), characterized by high initial values of ϕ , and issuing a capital structure of tranches against them. As a result of the tranching, the risk of default has been further concentrated into adverse economic states substantially increasing the appropriate risk compensation relative to rating-matched bonds. If investors fail to appreciate the impact of this mechanism on risk and continue to rely on their credit rating-to-price mapping deduced from single-name corporate bond markets, this new generation of securities will be significantly mispriced. In particular, investors willing to sell default protection on the highest rated tranches will be offered a yield insufficient to compensate them for the systematic risk they are asked to bear. This can be seen vividly in Table 6, where CDX tranches and single-name CDSs with similar loss rates, tend to trade at similar yields, despite their highly dissimilar economic risks,

and is equal to the yield spread that would arise if payoffs were discounted at the riskless rate.

³¹Elton 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 credit risk ratios that are twice as large for A-rated bonds than for BBB-rated bonds between 1996 and 2004.

as evidenced by vastly different ϕ values. This suggests that the investors in senior tranches either do not appreciate the risks that they are being asked to bear, or are severely constrained by the institutional requirement to hold highly-rated securities.

Figure 8 displays a scatter plot of the time series average value of $\log \phi$ for various securities against their model-implied loss rate, computed using their actual market yield spreads. The scattering of points sketch out a fairly smooth curve, which essentially represents the credit market’s pricing function based on credit rating, proxied here by loss rate. It appears that the non-equity CDX tranches are on average being priced to have yield spreads in line with rating-matched alternatives. In fact, these tranches offer spreads that are slightly higher than rating-matched alternatives, and thus appear to be priced attractively to an investor using this benchmark. Figure 8 also plots the average $\log \phi$ for the cheapest to supply security at each loss rate (i.e. the yield spreads associated with the digital market call options from Proposition 2). This curve represents a theoretical upper bound for $\log \phi$ at each loss rate. Interestingly, the average values of $\log \phi$ for the non-equity CDX tranches, computed using their *model* yield spreads, lie very close to this upper bound. This not only summarizes the earlier findings that the model yield spreads were much larger than the actual yield spreads for these securities, but suggests that these tranches were nearly optimally designed to be the cheapest securities to supply within their rating category. Moreover, since the values of the tranches sum to the value of the underlying portfolio, the overpricing of the non-equity tranches implies an underpricing of the first loss tranche. This can be seen in Figure 5, where the model yield spread is always below the actual yield spread for the 0%-3% tranche; and in Figure 8, where the mean actual ϕ for the 0%-3% tranche lies far above the model ϕ , and even above our theoretical upper bound.

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 (Peter 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, 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.

Finally, although the empirical focus of this paper is on CDOs comprised of investment grade credit default swaps, the underlying logic applies equally to any structured finance product consisting of assets with positive exposures to economic states. Structured products involving mortgages (CMOs) and corporate loans (CLOs) are likely to share these characteristics. Ultimately, how much of the repricing that has recently occurred in these markets is accounted for by the considerations of this paper depends on the relative importance of errors made in assessing the unconditional default

probabilities (i.e. credit ratings) of the underlying assets *vis a vis* their economic risk exposures.

7 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. Importantly, we highlight that the information credit rating agencies provide to their customers is inadequate for pricing. Forecasts of unconditional cash flows (i.e. credit ratings) 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 the 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 four to five times more risk 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.

A Proofs of Propositions

Proof of Proposition 1: Let $p_D^{X(p^*,N)}(s)$ denote the state-contingent tranche default probability, for a tranche whose unconditional default probability is p^* , and underlying asset pool is comprised of N bonds. Since we are considering a series of tranches with a fixed unconditional probability of default – under suitable regularity conditions allowing for interchange of integration and differentiation – we have:

$$\int_{s \leq \hat{s}} \frac{\partial p_D^{X(p^*,N)}}{\partial N}(s) \cdot \pi(s) ds + \int_{s > \hat{s}} \frac{\partial p_D^{X(p^*,N)}}{\partial N}(s) \cdot \pi(s) ds = 0 \quad (\text{A1})$$

where \hat{s} is chosen such that the derivative inside the integral is positive (negative) for states worse (better) than \hat{s} . Noting that the expected state-contingent tranche payoff is given by $1 - p_D^{X(p^*,N)}(s)$, we can differentiate the pricing equation for the digital tranche, under the same regularity conditions imposed above, to obtain,

$$\left. \frac{\partial B_\tau^X}{\partial N} \right|_{p_D^X = p^*} = - \int_{s \leq \hat{s}} \frac{\partial p_D^{X(p^*,N)}}{\partial N}(s) \cdot \left(\frac{q(s)}{\pi(s)} \right) \cdot \pi(s) ds - \int_{s > \hat{s}} \frac{\partial p_D^{X(p^*,N)}}{\partial N}(s) \cdot \left(\frac{q(s)}{\pi(s)} \right) \cdot \pi(s) ds \quad (\text{A2})$$

However, since s provides an ordering of economic states from worst to best, the value of $\frac{q(s)}{\pi(s)}$ – i.e. the price of receiving a dollar in state s per unit of probability of observing that state – will be a monotonically declining function of s , so long as the marginal investor is risk averse. Using this property and the mean value theorem, we know there will be two values \bar{m} and \underline{m} , with $\bar{m} > \underline{m}$, for which the above equation can be re-expressed as,

$$\left. \frac{\partial B_\tau^X}{\partial N} \right|_{p_D^X = p^*} = -\bar{m} \cdot \int_{s \leq \hat{s}} \frac{\partial p_D^{X(p^*,N)}}{\partial N}(s) \cdot \pi(s) ds - \underline{m} \cdot \int_{s > \hat{s}} \frac{\partial p_D^{X(p^*,N)}}{\partial N}(s) \cdot \pi(s) ds \quad (\text{A3})$$

Finally, from (A1), we know that the two integrals are equal in absolute value, with the first being positive and the second negative. If we denote the value of the first integral, η , we immediately have:

$$\left. \frac{\partial B_\tau^X}{\partial N} \right|_{p_D^X = p^*} = -(\bar{m} - \underline{m}) \cdot \eta < 0 \quad (\text{A4})$$

Consequently, increasing the number of securities in the underlying asset pool, N , while adjusting the tranche attachment point $X(p^*, N)$, such that the unconditional tranche default probability is fixed, causes the value of the digital tranche to decline monotonically. ■

Proof of Proposition 2: Consider the set of discount bonds with state-contingent default probabilities, $p_D(s)$, and unconditional default probability of p^* . The unconditional expected payoff for all such securities is, by the assumption of zero recovery in default, equal to $1 - p^*$. If we denote the time to maturity by τ , and the corresponding riskless rate of return by, r_f , the price of each

bond can be computed from,

$$\begin{aligned}
B_\tau(p^*) &= \exp(-r_f \cdot \tau) - \int_s p_D(s) \cdot \pi(s) \cdot \Lambda(s) ds \\
&= \exp(-r_f \cdot \tau) \left[1 - p^* \cdot \int_s \frac{p_D(s) \cdot \pi(s)}{\int_s p_D(s) \cdot \pi(s) ds} \cdot (\exp(r_f \cdot \tau) \cdot \Lambda(s)) ds \right] \\
&= \exp(-r_f \cdot \tau) [1 - p^* \cdot \bar{\Lambda}(p^*, p_D(s))]
\end{aligned} \tag{A5}$$

where we have taken advantage of the fact that the state prices themselves sum to the price of a τ -period riskless bond, $\exp(-r_f \cdot \tau)$, and defined $\Lambda(s)$ to be equal to the ratio of the state's price to its probability, $\frac{q(s)}{\pi(s)}$. As in the previous proof, note that if the marginal investor is risk averse, $\Lambda(s)$ will be monotonically decreasing in s . Finally, since $\bar{\Lambda}(p^*, p_D(s))$, reflects the average value of $\Lambda(s)$ in the states in which the bond defaults, it follows that the bond with the lowest price (i.e. largest yield spread), fails to pay on a set with measure p^* containing the worst economic states. ■

Proof of Proposition 3: Consider a sequence of digital tranches, whose attachments points $X(p^*, N)$, are set as a function of the number of securities in the underlying portfolio, N , to maintain the tranche's unconditional default probability fixed at p^* . As $N \rightarrow \infty$, the state-contingent tranche payoff converges in probability, as follows,

$$\lim_{N \rightarrow \infty} \Phi \left(\sqrt{N} \cdot \frac{(X(p^*, N) - p_D(s))}{\sqrt{p_D(s) \cdot (1 - p_D(s))}} \right) = \begin{cases} 0 & s < \tilde{s} \\ 1 & s \geq \tilde{s} \end{cases} \tag{A6}$$

where \tilde{s} is the p^* -the percentile of the τ -period distribution of economic states, guaranteeing that the unconditional tranche default probability is p^* . If the cumulative distribution of economic states is denoted by $\Pi(s)$, then the attachment point of the asymptotically diversified tranche can be recovered from, $X(p^*, \infty) = p_D(\Pi^{-1}(p^*))$. Since the asymptotically diversified tranche pays zero on a set with measure p^* containing the worst economic states, and one elsewhere, its payoff converges to the payoff of the cheapest asset to supply with that level of default risk. ■

B Computing State Prices

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

$$\sigma(x, \tau) = a + b \cdot \tanh(-c \cdot \ln x) \tag{B1}$$

$$\sigma(x, \tau) = a + b \cdot \exp(-c \cdot x) \tag{B2}$$

where x is the option moneyness, defined as the ratio of the option strike price to the prevailing τ -period futures 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. We require that $\sigma(x, \tau)$ converge to one half of its at-the-money value, $\sigma(1, \tau)$, as $x \rightarrow \infty$. Finally, in order to preclude arbitrage opportunities, we require that the Arrow-Debreu prices implied by our implied volatility function be strictly positive for all moneyness values.

For the hyperbolic tangent specifications, the first and second derivatives of implied volatility function with respect to the option moneyness are given by:

$$\frac{d\sigma(x, \tau)}{dx} = -\left(\frac{b \cdot c}{x}\right) \cdot \operatorname{sech}^2(c \cdot \ln x) \quad (\text{B3})$$

$$\frac{d^2\sigma(x, \tau)}{dx^2} = \left(\frac{b \cdot c}{x^2}\right) \cdot \operatorname{sech}^3(c \cdot \ln x) \cdot (\cosh(c \cdot \ln x) + 2c \cdot \sinh(c \cdot \ln x)) \quad (\text{B4})$$

The corresponding derivatives for the exponential implied volatility function are:

$$\frac{d\sigma(x, \tau)}{dx} = -b \cdot c \cdot \exp(-c \cdot x) \quad (\text{B5})$$

$$\frac{d^2\sigma(x, \tau)}{dx^2} = b \cdot c^2 \cdot \exp(-c \cdot x) \quad (\text{B6})$$

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. Denoting the τ -period futures price at time t by $F_{t,\tau}$, we have,

$$\left. \frac{\partial^2 C^{BS}}{\partial K^2} \right|_{K=x \cdot F_{t,\tau}} = e^{-r_f \cdot \tau} \cdot \frac{\phi(d_2(x))}{x \cdot F_{t,\tau} \cdot \sigma(x, \tau) \cdot \sqrt{\tau}} \quad (\text{B7})$$

$$\left. \frac{\partial C^{BS}}{\partial \sigma} \right|_{K=x \cdot F_{t,\tau}} = e^{-r_f \cdot \tau} \cdot x \cdot F_{t,\tau} \cdot \phi(d_2(x)) \cdot \sqrt{\tau} \quad (\text{B8})$$

$$\left. \frac{\partial^2 C^{BS}}{\partial \sigma^2} \right|_{K=x \cdot F_{t,\tau}} = e^{-r_f \cdot \tau} \cdot x \cdot F_{t,\tau} \cdot \phi(d_2(x)) \cdot \sqrt{\tau} \cdot \frac{d_1(x) \cdot d_2(x)}{\sigma(x, \tau)} \quad (\text{B9})$$

$$\left. \frac{\partial^2 C^{BS}}{\partial K \partial \sigma} \right|_{K=x \cdot F_{t,\tau}} = e^{-r_f \cdot \tau} \cdot \frac{d_1(x) \cdot \phi(d_2(x))}{\sigma(x, \tau)} \quad (\text{B10})$$

where:

$$d_1(x) = -\frac{\ln x}{\sigma(x, \tau) \cdot \sqrt{\tau}} + \frac{1}{2} \cdot \sigma(x, \tau) \cdot \sqrt{\tau} \quad (\text{B11})$$

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

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.

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 and Vellore M. Kishore, 1996, "Almost Everything You Wanted to Know About Recoveries on Defaulted Bonds," *Financial Analysts Journal* 52, 57-64.
- [4] Arrow, Kenneth J., 1964, "The Role of Securities in the Optimal Allocation of Risk Bearing," *Review of Economic Studies* 31, 91-96.
- [5] Azizpour, Shahriar, and Kay Giesecke, 2008, "Premia for Correlated Default Risk," Stanford University working paper.
- [6] Bank for International Settlements, 2005, "The Role of Ratings in Structured Finance: Issues and implications," (available from: <http://www.bis.org>).
- [7] Bates, David, 2006, "The Market Price of Crash Risk," University of Iowa working paper.
- [8] Berd, Arthur, Sunita Ganapati, Roy Mashal, Marco Naldi, Dominic O'Kane and Claus Pedersen, 2003, *The Lehman Brothers Guide to Exotic Credit Derivatives*, London: Risk Waters Group.
- [9] 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.
- [10] Black, Fischer and John C. Cox, 1976, "Valuing Corporate Securities: Some Effects of Bond Indenture Provisions," *Journal of Finance* 31, 351-367.
- [11] Black and Myron Scholes, 1973, "The pricing of options and corporate liabilities," *Journal of Political Economy* 81, 617-654.
- [12] Bliss, Robert and Nikolaos Panigirtzoglou, 2004, "Option-Implied Risk Aversion Estimates," *Journal of Finance* 59, 407-446.
- [13] Bondarenko, Oleg, 2003, "Why Are Put Options So Expensive?" UIC working paper.
- [14] Breeden, Douglas and Robert Litzenberger, 1978, "Prices of state-contingent claims implicit in option prices," *Journal of Business*, 621-651.
- [15] Broadie, Mark, Chernov, Mikhail, and Michael Johannes, 2008, "Understanding Index Option Returns," Columbia University working paper.
- [16] Brown, Greg, and Klaus B. Toft, 1999, "Constructing binomial trees from multiple implied probability distributions," *Journal of Derivatives*, 7, 83-100.
- [17] 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.

- [18] Campa, Jose M., Chang, Kevin, and Robert L.Reider, 1997, "ERM bandwidth for EMU and after: Evidence from foreign exchange options," *Economic Policy* 24, 55–89.
- [19] Cantor, Richard, David T. Hamilton, Sharon Ou, and Praveen Varma, 2005, *Default and Recovery Rates of Corporate Bond Issuers, 1920-2005*, Moody's Investor Service: Global Credit Research, 2005.
- [20] Collin-Dufresne, Pierre, Robert S. Goldstein, and J. Spencer Martin, 2001, "The Determinants of Credit Spread Changes," *Journal of Finance* 56, 2177-2208.
- [21] 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).
- [22] Debreu, Gerard, 1959, *Theory of Value*, New York: Wiley.
- [23] Delianedis, Gordon, and Robert Geske, 2001, "The Components of Corporate Credit Spreads: Default, Recovery, Tax, Jumps, Liquidity, and Market Factors," UCLA working paper.
- [24] DeMarzo, Peter, 2005, "The pooling and tranching of securities: A model of informed intermediation," *Review of Financial Studies*, 1-35.
- [25] Driessen, Joost, 2005, "Is Default Event Risk Priced in Corporate Bonds?" *Review of Financial Studies*, 18, p. 165-195.
- [26] Duffie, Darrell, and Gârleanu, 2001, "Risk and Valuation of Collateralized Debt Obligations," *Financial Analysts Journal*, 41-59.
- [27] Duffie and Kenneth J. Singleton, 1999, "Modeling Term Structures of Defaultable Bonds," *Review of Financial Studies* 12, 687-720.
- [28] Duffie and Singleton, 2003, *Credit Risk*, Princeton University Press.
- [29] Eckner, Andreas, 2008, "Risk Premia in Structured Credit Derivatives," Stanford University working paper.
- [30] 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.
- [31] Eom, Young Ho, Jean Helwege, and Jing-Zhi Huang, 2004, "Structural Models of Corporate Bond Pricing," *Review of Financial Studies*, 17, p. 499-544.
- [32] Fama, Eugene F., 1970, "Efficient Capital Markets: A Review of Theory and Empirical Work," *Journal of Finance* 25, 383-417.
- [33] Fama and Kenneth R. French, 1993, "Common Risk Factors in the Returns on Stocks and Bonds," *Journal of Financial Economics*, 33, 3-56.
- [34] 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.
- [35] Hull, John, Mirela Predescu, and Alan White, 2005, "Bond prices, default probabilities, and risk premiums," *Journal of Credit Risk*, 53-60.

- [36] Hull, Predescu, and White, 2006, “The Valuation of Correlation-dependent Credit Derivatives Using a Structural Model,” University of Toronto working paper.
- [37] Jackwerth, Jens C., “Option Implied Risk-Neutral Distributions and Implied Binomial Trees: A Literature Review,” *Journal of Derivatives* 7, 66-82.
- [38] 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
- [39] Kakodkar, Atish and Barnaby Martin, 2004, “The Standardized Tranche Market Has Come Out of Hiding,” *Journal of Structured Finance*, Fall, 76-81.
- [40] Leland, Hayne, 1994, “Corporate Debt Value, Bond Covenants, and Optimal Capital Structure,” *Journal of Finance* 49, 1213-1252.
- [41] Lintner, John, 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.
- [42] 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.
- [43] Longstaff and Arvind Rajan, 2007, “An Empirical Analysis of the Pricing of Collateralized Debt Obligations,” *Journal of Finance*, forthcoming.
- [44] Merton, Robert C., 1973, “Theory of rational option pricing,” *Bell Journal of Economics and Management Science*, 141-183.
- [45] Merton, 1974, “On the pricing of corporate debt: The risk structure of interest rates,” *Journal of Finance*, 449-470.
- [46] Pan, Jun, and Singleton, 2006, “Default and recovery implicit in the term structure of sovereign CDS spreads,” MIT working paper.
- [47] Rosenberg, Joshua and Robert F. Engle, 2002, “Empirical pricing kernels,” *Journal of Financial Economics*, 64, 341-372.
- [48] Schaefer, Stephen and Ilya Strebulaev, 2005, “Structural Models of Credit Risk are Useful: Evidence from Hedge Ratios on Corporate Bonds,” Stanford GSB working paper.
- [49] Schönbucher, Philipp, 2000, “Factor Models for Portfolio Credit Risk,” Bonn University working paper.
- [50] Shimko, David, 1993, “Bounds of probability,” *Risk* 6, 33–37.
- [51] Sharpe, William F., 1964, “Capital asset prices: A theory of market equilibrium under conditions of risk,” *Journal of Finance* 19, 425-442.
- [52] Vasicek, Oldrich, 1987, “Probability of Loss on Loan Portfolio,” KMV Corporation.
- [53] Vasicek, 1991, “Limiting Loan Loss Distribution,” KMV Corporation.
- [54] 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.

- [55] Zhou, Chunsheng, 2001, "An Analysis of Default Correlations and Multiple Defaults," *Review of Financial Studies* 14, 555-576.