

NBER WORKING PAPER SERIES

ON THE RELATIVE PRICING OF LONG MATURITY S&P 500 INDEX OPTIONS
AND CDX TRANCHES

Pierre Collin-Dufresne
Robert S. Goldstein
Fan Yang

Working Paper 15734
<http://www.nber.org/papers/w15734>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
February 2010

The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

© 2010 by Pierre Collin-Dufresne, Robert S. Goldstein, and Fan Yang. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

On the Relative Pricing of long Maturity S&P 500 Index Options and CDX Tranches
Pierre Collin-Dufresne, Robert S. Goldstein, and Fan Yang
NBER Working Paper No. 15734
February 2010
JEL No. G12,G13

ABSTRACT

We investigate a structural model of market and firm-level dynamics in order to jointly price long-dated S&P 500 options and tranche spreads on the five-year CDX index. We demonstrate the importance of calibrating the model to match the entire term structure of CDX index spreads because it contains pertinent information regarding the timing of expected defaults and the specification of idiosyncratic dynamics. Our model matches the time series of tranche spreads well, both before and during the financial crisis, thus offering a resolution to the puzzle reported by Coval, Jurek and Stafford (2009).

Pierre Collin-Dufresne
Graduate School of Business
Columbia University
Uris Hall 404
3022 Broadway
New York, NY 10027
and NBER
pc2415@columbia.edu

Fan Yang
Wpłxgtuk{ "qh'O kppguvc
321 19th Cve Uouth
yang0946@umn.edu

Robert S. Goldstein
University of Minnesota
Finance Department
3-125 Carlson School of Management
321 19th Avenue South
Minneapolis, MN 55455
and NBER
golds144@umn.edu

1 Introduction

The explosive growth of the credit derivative market, and the ensuing economic crisis, have generated much interest in the pricing and hedging of credit derivatives. This market is composed of both single-name products, such as the credit default swap (CDS), and multi-name products such as the CDX index basket default swap and tranches based on this index. While the spread on the CDX index can be mostly determined by observing the CDS spreads on the 125 underlying firms that compose the index, the pricing of CDX tranche spreads depends crucially on how one models default correlation. That is, whereas the pricing of the CDX index depends only on the marginal default probabilities of the individual firms, the pricing of the tranches depends on their joint probability distribution.

One framework often used for determining spreads on individual corporate bonds is the so-called structural model of default (Merton (1974)), which is based on the insights of Black and Scholes (1973). Risk neutral dynamics are specified for firm value, and default is assumed to occur at maturity if firm value falls below a default boundary, which in the original model is specified to equal the face value of debt outstanding. One interesting prediction of this framework is that spreads depend only on the firm's total variance, and not on the fraction of variance due to market risk and idiosyncratic risk.¹ This is because only the marginal default probability is needed to determine bond or CDS spreads.

In contrast, because CDX tranche spreads depend on correlation, their prices are rather sensitive to the composition of market and idiosyncratic risks of the underlying firms. One of the first frameworks for modeling default correlation were the so-called copula models (Vasicek (1987), Li (2000)). These models are effectively structural models of default embedded within a CAPM framework to model default correlation. In particular, the standard copula model can be interpreted as modeling returns as a sum of two factors: a market factor and an idiosyncratic factor. Just as in the standard structural model, default occurs if the the sum of these two factors drops below a specified default boundary. The level of correlation is driven by the relative weights specified on the market and idiosyncratic components. Unfortunately, the original copula framework is not dynamically consistent. For example, when pricing tranches with different maturities, the standard copula approach has been to specify a different default boundary location and a different market beta for each maturity, and then assume that default

¹Note that if we condition on the P -measure default probability instead of the initial leverage ratio, then the fraction of risk due to systematic risk would matter. Indeed, consider two firms that have the same total risk, default probability and recovery rate. The firm with higher beta (and hence higher expected return) will need to start with higher leverage to have the same default probability as the low-beta firm. As such, the higher beta-firm will command a higher spread (Chen, Collin-Dufresne and Goldstein (2008)).

can occur only at the maturity of the security in question. Moreover, copula models are inherently static, and thus cannot be used, for example, to price options on CDO's (Duffie (2004)).

In the original Gaussian copula models, both market and idiosyncratic factors were assumed to be normally distributed. Just as it has become standard to quote option prices in terms of its Black/Scholes implied volatility (in spite of the well-documented flaws in the model's predictions), it has become standard to quote correlations implied from the Gaussian copula. Many improvements to the Gaussian copula have been suggested to improve upon its empirical failures.²

In a recent influential paper, Coval, Jurek and Stafford (CJS, 2009) investigate the pricing of CDO tranches within a copula framework by using S&P 500 option prices to identify the market component of returns. Their approach is motivated by Breeden and Litzenberger (1978), who show how state prices can be extracted from option prices.³ Given state prices, CJS then calibrate the idiosyncratic component of returns from observed equity returns, and the location of the default boundary so that their model matches perfectly the five-year CDX index (i.e., the risk-neutral expected loss on the index). They find that observed spreads on the equity tranche are too high, and observed spreads on more senior tranches are too low relative to those predicted by their model.⁴ Since their results seem robust along many dimensions, CJS conclude that sellers of senior protection were writing insurance contracts on "economic catastrophe bonds" without realizing the magnitude of systematic risk they were exposed to, and thus not demanding adequate compensation for that risk. They conjecture that agents had purchased senior claims solely with regard to their credit rating (which reflects their P -measure default probability), but ignored the large systematic risk associated with holding a senior claim written on an underlying diversified portfolio. Their findings are puzzling, since traders in the CDX market are typically thought of as being rather sophisticated. Thus it would be surprising to find them accepting so much risk without fair compensation.

In this paper, we also investigate the relative prices of S&P 500 index options and spreads on the CDX index and tranches. An important difference between our framework and that of CJS (and most other copula models) is that we specify a fully dynamic model that allows

²See, for example, Andersen and Sidenius (2004), Hull and White (2004), Duffie and Garleanu (2001).

³Since it is well-known that fat tails in the common density are needed to fit tranche prices. CJS propose an 'out-of-sample' calibration of the common market factor distribution. Other approaches to match fat tails includes Hull and White (2004), who use a double t-distribution.

⁴For their benchmark case, spreads are approximately two-times too small for the 3-7% tranche, four-times too small for the 7-10% tranche, five-times too small for the 10-15% tranche, and three-times too small for the 15-30% tranche.

us to jointly and consistently price these securities across the maturity spectrum. Following Black and Cox (1976), we specify the default event as the first time firm value drops below the default boundary, instead of limiting default to occur only at maturity. This allows us to account for the possibility of early default, which can significantly impact the cash flows of protection buyers and sellers of CDX tranches.

While there are many differences between our framework and theirs, the most crucial distinction that drives our results is that CJS calibrate their model to match only the five-year CDX index spread, whereas we calibrate our model to match the entire term structure of CDX index spreads – that is, spreads across *all* maturities. Calibrating to shorter horizon CDX index spreads is important because they contain pertinent information regarding the timing of expected defaults and the specification of idiosyncratic dynamics. Using this calibration approach, we match the time series of tranche spreads well, both during and prior to the crisis period, thus providing a resolution to the puzzle of CJS.

To gain some intuition for why it is essential to calibrate the model to match the term structure of CDX index levels, we recall a well-documented failure of (diffusion-based) structural models of default: they dramatically underpredict default rates for investment grade debt at short maturities. This is true both under the risk-neutral measure, where structural models predict negligible credit spreads at short maturity (Jones, Mason and Rosenfeld (1984), Eom, Helwege Huang (2005)), and under the true measure (Leland (2004)). To match short-term spreads on investment grade debt, it is necessary to include idiosyncratic jumps into firm dynamics. (e.g., Zhou (2001)). We calibrate idiosyncratic jump size and intensities to match observed CDX index spreads at short horizons. In sum then, whereas S&P 500 option prices are useful for identifying market dynamics, the term structure of CDX index levels is useful for identifying idiosyncratic dynamics.

Calibrating the idiosyncratic component of return dynamics to the term structure of CDX index spreads impacts tranche spreads in two important ways. First, it increases (risk neutral) expected losses (i.e., defaults) at shorter horizons. Without this calibration approach, defaults are *backloaded*, that is, occur later than actual market expectations. The timing of defaults is especially crucial for the equity tranche, where the buyer of protection gets to reduce her payments as soon as the first default occurs. This helps explain why the implied equity tranche spread in the CJS framework tends to be too low relative to the market quote. Moreover, since their model matches the five year CDX index (and therefore, five year risk neutral expected losses), the downward bias on the equity tranche spread automatically generates an upward bias on the spreads of the senior tranches. It is interesting to note that, in practice, the timing

of defaults is understood to be so crucial for the determination of the equity tranche spread that it is common to contractually specify its cash flows in a different manner (i.e., with an up-front premium) compared to the other tranches (i.e., full-running premium) so as to reduce its sensitivity to the timing of defaults.⁵ As we demonstrate below, our model matches observed equity tranche spreads well for both forms of contractual payments.

The second important way our calibration approach impacts tranche spreads is by increasing the proportion of risk that is idiosyncratic compared to systematic. The intuition for how this impacts tranche spreads stems from the central limit theorem: if in fact all risk were idiosyncratic, then due to the large number (namely, 125) of firms in the index, the distribution of fractional losses would be tightly peaked around expected losses (which were below the 3% equity loss detachment point during the period studied by CJS). As one increases the proportion of systematic risk, the loss distribution widens because losses become more dependent on market performance. Below we show that if the model is not calibrated to match the term structure of CDX index spreads, then there is an insufficient level of idiosyncratic risk specified, and the loss distribution is not sufficiently peaked. This forces too much of the loss distribution (compared to the market's expectations) to fall above the 3% equity loss detachment point, implying too high of a probability where sellers of protection of the more senior tranches will have to pay out, in turn biasing upward tranche spreads for the more senior tranches. Again, since CJS calibrate their model to match the 5 year CDX index, the upward bias in senior tranche spreads automatically generates a downward bias in equity tranche spreads.

One more issue that arises in our investigation is that we show one cannot uniquely determine market dynamics from option contracts that are actually traded. This is because the strike prices of traded options do not span far enough in the moneyness dimension to identify the (risk neutral) probabilities of catastrophic crashes. As such, there are many choices of market dynamics that can be specified that generate nearly identical prices for those options with strikes that are actually traded, but generate very different prices for the super-senior tranche. Therefore, we calibrate our model to match the super-senior tranche spread, and price all other tranche spreads out-of-sample.⁶

While CJS calibrate state prices only for a maturity of five years, our framework requires a

⁵As we discuss in equations (14)-(15) below, instead of setting the equity tranche premium so that the present value of the premium equals the present value of the protection, often there is also an up-front premium paid to sellers of equity tranche protection.

⁶We are confident we could match tranche spreads even better than what is shown below if we were to also calibrate market dynamics to match other tranche spreads. But such an additional exercise would mostly mimic the contributions inherent in papers such as Mortensen (2006), Longstaff and Rajan (2008), and Eckner (2009), who have already demonstrated the ability of a few in-sample state variables to match the cross section of tranche spreads well.

dynamic model capable of producing state prices for all maturities. As documented previously in the literature, it is not a trivial task to specify market dynamics to match long maturity options, as observed option implied volatilities fail to flatten out as maturity increases, in contrast to the predictions of most models.⁷ To circumvent this problem, we allow the dividend yield to be stochastic. It turns out that this can have a significant impact on the implied volatility surface, especially if the parameterization implies non-stationary dynamics under the risk neutral measure.⁸ While non-stationarity may at first seem undesirable, it has a theoretical justification. Indeed, we show in the Appendix that, in the habit formation model of Campbell and Cochrane (1999), the endogenously determined dividend yield is a function of their ‘consumption surplus’ habit variable and that this variable exhibits non-stationarity under the risk neutral measure (even though it is stationary under the historical measure). Moreover, this feature is responsible for their model’s ability to capture long term predictability of market returns.

There is a large and increasing literature on correlated defaults.⁹ Our contribution with respect to the literature is to investigate the relative pricing across the stock option and CDO markets.¹⁰ Further, our fully dynamic and self-consistent framework allows us to investigate the sensitivity of standard approximations made in the literature. For example, we find that accounting for firm heterogeneity, dynamic capital structure, ‘industry correlations’, or stochastic interest rates has only a secondary impact on tranche pricing. In contrast, calibrating to the term structure of CDX spreads is essential for jointly modeling S&P 500 option prices and CDX tranche spreads.

The rest of the paper is as follows: In Section 2 we propose a joint model for equity index options and CDO tranches. In Section 3 we discuss the data and our calibration approach. In Section 4 we report the results for option prices as well as the cross sectional implications for tranche prices of our various model specifications. In Section 5 we investigate the time series properties of our model. In Section 6, we investigate the robustness of our results. We conclude in Section 7.

⁷See Backus, Foresi and Wu (1997), Das and Sundaram (1999), Carr and Wu (2003).

⁸With that said, our results for tranche spreads are not driven by our choice of modeling market returns, as we demonstrate in the robustness section.

⁹Examples include Duffie and Garleanu (2001), Collin-Dufresne, Goldstein and Helwege (2003), Hull and White (2004), Longstaff, Mithal and Neis (2005), Mortensen (2006), Longstaff and Rajan (2008), Giesecke and Goldberg (2005), Bakshi, Madan and Zhang (2006), Das et al (2006), Das et al (2007), Duffie, Saita and Wang (2007), Jorion and Zhang (2009)

¹⁰Cremers, Driessen and Maenhout (2008) show option prices on individual firms are consistent with prices from the CDS market.

2 A joint structural model for equity index options and CDO tranches

This paper investigates the relative pricing of CDX tranche spreads and S&P 500 option prices within a structural framework. In order to do so, we need a model of market returns and individual firm level returns. In this section, we introduce an affine option pricing model (SVDCJ) for market returns that extends the standard SVCJ model of Broadie, Chernov and Johannes (2009) to better account for some features inherent in the data of long-maturity option prices. We then introduce a CAPM-like structural model for firm level returns that allows us to price basket CDS and CDO tranches.

2.1 Market dynamics for long maturity option prices

There is a long tradition, since Breeden and Litzenberger (1978), to extract implied state price densities from quoted option prices. One common approach, followed by CJS, is to use a ‘local volatility’ model, which specifies a flexible form for the implied volatility function which is then calibrated to match observed option prices across strikes and maturity.¹¹ This approach is particularly well suited when the amount of data is large, and tends to provide accurate estimates for the implied density if the range of interpolation/extrapolation is not too far from strikes for which option prices are available. Unfortunately, to study tranche prices, and especially the super-senior tranche, we will have to ‘extrapolate’ the density to regions that are far from any strikes of quoted option prices. Therefore, we prefer to use a fully dynamic model where return dynamics under the risk neutral measure are specified. This approach guarantees that our ‘extrapolated’ volatility surface is both consistent (i.e., arbitrage-free) and generated from a model with economically motivated dynamics. One nice feature (and necessary for our needs) of having a fully dynamic option pricing model is that, once calibrated, we can obtain the state price density for all strikes and all maturities. In contrast, CJS focus on state prices for the 5-year maturity only, and specify a model where this information is sufficient to price tranche spreads.

There is a large literature on testing parametric option pricing models. Early papers that test various specifications of option pricing models include Bates (2000), Pan (2002), Eraker, Johannes, Polson (2003), Eraker (2004), Anderson and Andreasen (2000).¹² The recent

¹¹See, for example, Derman and Kani (1994), Dupire (1994) Rubinstein (1994), Dumas, Fleming and Whaley (1998).

¹²More recently, researchers have investigated whether option prices are consistent with asset prices in different markets. For example, Cremers, Driessen and Maenhout (2008) demonstrate that the implied volatility smirk of option prices on individual stocks are mostly consistent with credit spreads on the same firm.

paper by Broadie, Chernov and Johannes (BCJ 2009) presents an encompassing test of various specifications proposed in the literature using an extensive data set on S&P 500 futures options. Their preferred model is the so-called SVCJ model which allows for stochastic volatility and correlated jumps in both stock returns and volatility processes. We therefore choose to build on their preferred model, which we extend by introducing a stochastic dividend yield. BCJ and indeed most of the literature has focused on options with relatively short maturity (i.e., less than 6 months). However, since our focus is the pricing of 5-year CDX tranches, it is necessary to back out long-maturity state prices, implying the need to look at long-maturity option prices. The maturity of exchange-traded options on the S&P 500 is typically limited to under three years, but on the OTC market it is possible to obtain longer dated options. CJS calibrate their model to five year maturity quotes obtained from Citigroup. We perform our initial calibration by matching to the same prices as reported in CJS.

While the so-called SVCJ jump-diffusion model performs very well at capturing time-series and cross-sectional properties of option returns, it performs better for short-dated options. For longer maturity options the model-predicted implied volatility surface flattens (see Figure 1), in contradiction with the longer-term maturity option data available. This surprising pattern found in long maturity option prices, that implied volatilities, expressed as a function of a standard measure of moneyness, do not flatten out, implies a violation of the central limit theorem's underlying assumptions (Carr and Wu (2003)). Indeed, the slope and curvature of the volatility 'smirk' are mostly controlled by the skewness and kurtosis of the distribution of market prices. But if the central limit theorem holds, it predicts that the distribution of log-returns should converge to normal over sufficiently long maturity.¹³ In order to overcome the central limit theorem's implications, Carr and Wu (2003) propose a log-stable process with infinite second moments. While their model matches option prices well, one unattractive feature of their model is that the variance of log returns is infinite in their model (despite the fact that all moments of the spot price and therefore option prices are finite). This could lead to infinite prices of certain volatility derivatives for example. Instead, we propose an alternative mechanism to fit longer maturity implied volatility skews.

While there are many possible choices (e.g., state variables driving interest rates, volatility, intensity, etc.), here we consider a state variable that has been mostly ignored in the options literature, even though it is fundamental in the returns literature: the dividend yield (i.e., the inverse of the price/dividend ratio that is the focus of many macro-finance papers such as Campbell and Cochrane (1999) or Bansal and Yaron (2003)). As motivation, we note that

¹³See Das and Sundaram (1999), Backus, Foresi and Wu (2004).

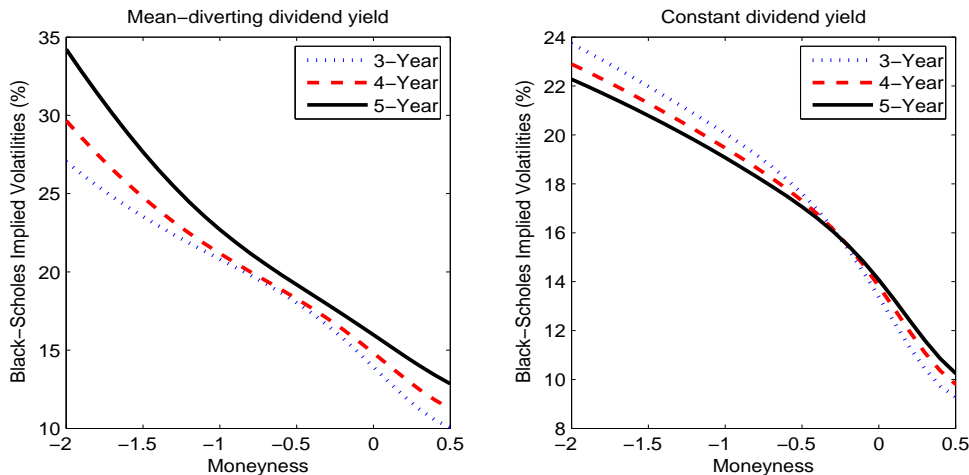


Figure 1: Implied volatility curves (IVF) as a function of moneyness for different maturities for two different models: our non-stationary dividend yield model, and constant dividend yield. We follow Carr and Wu(2003) definition of moneyness $\frac{\log(K/F)}{\sigma\sqrt{T}}$. Note how the five-year IVF does not flatten out compared to the three- and four-year for the non-stationary model, in contrast to the constant dividend model.

in the habit formation model of Campbell and Cochrane (1999), the price-dividend ratio is a function of a single state variable s , the so-called ‘surplus consumption ratio’, whose dynamics are stationary under the actual measure, but non-stationary under the risk-neutral measure. (See Appendix 1). Therefore we introduce a stochastic dividend yield, with possible non-stationary dynamics under the risk-neutral measure. Specifically, we define M_t as the value of the market portfolio, V_t the volatility of the market, and $\delta(t)$ the dividend yield of the market. We specify risk-neutral market dynamics as:

$$\frac{dM_t}{M_t} = (r - \delta) dt + \sqrt{V_t} dw_1^Q + (e^y - 1) dq - \bar{\mu}_y \lambda^Q dt - (e^{y_C} - 1) (dq_C - \lambda_C^Q dt) \quad (1)$$

$$dV_t = \kappa_V (\bar{V} - V_t) dt + \sigma_V \sqrt{V_t} (\rho dw_1^Q + \sqrt{1 - \rho^2} dw_2^Q) + y_V dq \quad (2)$$

$$d\delta_t = \kappa_\delta (\bar{\delta} - \delta_t) dt + \sigma_\delta \sqrt{V_t} (\rho_1 dw_1^Q + \rho_2 dw_2^Q + \sqrt{1 - \rho_1^2 - \rho_2^2} dw_3^Q) + y_\delta dq. \quad (3)$$

Here, dw_1^Q and dw_2^Q are Brownian motions, and dq is a jump process with a constant jump intensity λ^Q . The market, the variance and the dividend yield jump contemporaneously. The jump size of the volatility follows an exponential distribution, $y_V \sim \exp(1/\mu_V)$. The jump size in market is originally assumed to be normal conditional on the jump size of V_t , $y|y_V \sim N(\rho_q y_V + \mu_y, \sigma_y)$. The jump size in dividend yield is assumed to be independent

normal, $y_\delta \sim N(\mu_d, \sigma_d)$. The compensator for the jump in the market price is

$$\bar{\mu}_y = \mathbb{E}[e^y] - 1 = \frac{e^{\mu_y + \frac{1}{2}\sigma_y^2}}{1 - \rho_q \mu_v} - 1. \quad (4)$$

We note that this model is the standard SVCJ model studied in the literature (e.g., BCJ (2009)) extended for a stochastic dividend yield with a correlated jump. We therefore refer to this model as the SVDCJ model. We note that in addition we have explicitly added a jump q_C with deterministic jump size y_C in the market dynamics, even though it could be subsumed in the standard jump dq , because we want to emphasize the importance of a ‘catastrophic’ jumps for the pricing of super-senior tranches in our calibrations. The importance of catastrophic jumps for asset pricing has been emphasized in Rietz (1983) and more recently revisited in Barro (2008), Backus, Chernov and Martin (2009) and others. We demonstrate below that there are many calibrations of this model that would match well the option prices for those strikes that are actually traded. Therefore, in order to uniquely identify market dynamics, we calibrate the model to match the data on both the super-senior tranche spreads and S&P 500 options, leaving all other tranche spreads to be priced out-of-sample.

Given that the model is specified to have affine dynamics, it follows that the characteristic function is exponential affine (Duffie, Pan and Singleton (2000)). As such, European option prices can be solved numerically by applying the Fast Fourier Transformation(FFT) on the log market characteristic function (Heston (1996), Carr and Madan (1998)). The solution of the log market characteristic function is given in the Appendix 2.

2.2 Firm dynamics and structural default model

Given the market dynamics, we assume that individual firm dynamics are specified as:

$$\begin{aligned} \frac{dA}{A} + \delta_A dt - rdt &= \beta \left(\sqrt{V_t} dw_1^Q + (e^y - 1) dq - \bar{\mu}_y \lambda^Q dt \right) + (e^{y_C} - 1) (dq_C - \lambda_C^Q dt) \\ &\quad + \sigma_i dw_i + (e^{y_i} - 1) (dq_i - \lambda_i^Q dt). \end{aligned} \quad (5)$$

This is basically the standard ‘CAPM’ like equation for individual firm’s asset return, where β denotes the loading of each firm’s asset return dynamics on the market return. The only difference from CAPM is that we assume each firm has a loading of 1 on the catastrophic event. That is, unlike the standard market risk, we assume that all firms have the same exposure to ‘catastrophic events’. The idiosyncratic jump size and intensity are constants. Under these specifications, the log market and log asset have dynamics

$$d \log M_t = \left(r - \delta_t - \frac{1}{2} V_t - \bar{\mu}_y \lambda^Q - (e^{y_C} - 1) \lambda_C^Q \right) dt + \sqrt{V_t} dw_1^Q + y dq + y_C dq_C$$

$$\begin{aligned}
d \log A_t = & \left(r - \delta_A - \frac{1}{2} \beta^2 V_t - \frac{1}{2} \sigma_i^2 - \beta \bar{\mu}_y \lambda^Q - (e^{y_C} - 1) \lambda_C^Q - (e^{y_i} - 1) \lambda_i^Q \right) dt \\
& + \beta \sqrt{V_t} dw_1^Q + \sigma_i dw_i + \log [1 + \beta (e^y - 1)] dq + y_C dq_C + y_i dq_i.
\end{aligned} \tag{6}$$

Following Black and Cox (1976) and others, we specify that default occurs the first time firm value falls below a default threshold A_B . Therefore default arrival time for the typical firm i with asset dynamics $A_i(t)$ is defined as:

$$\tau_i = \inf\{t : A_i(t) \leq A_B\}. \tag{7}$$

We assume that upon default the debt-holder recover a fraction of the remaining asset value $(1 - \ell)A_B$ where ℓ is the loss rate.

2.3 Basket CDS index and CDO tranche spreads

We consider next the pricing of baskets of CDS and CDO tranches. Eventually, we want to use the model to price the DJ CDX North American Investment Grade Index and the tranches associated with it. This index is a basket CDS written on an equally weighted portfolio of 125 investment grade names. It can essentially be thought of as a portfolio of 125 liquid five-year credit default swaps (CDS) with investment grade status. The spread of a CDS on a firm is closely related to the credit spread of a bond on the same firm.¹⁴ The running spread on the CDX index is closely related to a weighted average of CDS spreads.¹⁵ To determine this spread, the present value of cash flows that go to the protection buyer and protection seller are set equal to each other. These cash flows are specified as

$$V_{idx,prem}(S) = S E \left[\sum_{m=1}^M e^{-rt_m} (1 - n(t_m)) \Delta + \int_{t_{m-1}}^{t_m} du e^{-ru} (u - t_{m-1}) dn_u \right] \tag{8}$$

$$V_{idx,prot} = E \left[\int_0^T e^{-rt} dL_t \right]. \tag{9}$$

Here, we have defined the number of defaults in the portfolio by $n(t) = \sum_i \mathbf{1}_{\{\tau_i \leq t\}}$, and the cumulative loss in the portfolio as:

$$L(t) = \sum_i \mathbf{1}_{\{\tau_i \leq t\}} R_i(\tau_i) \tag{10}$$

¹⁴Empirically, CDS spreads are close to credit spreads only when the risk-free benchmark is estimated to be well above Treasuries. In normal times, the swap/Libor curve is a better proxy than the Treasury curve. See, for example, Blanco Brennan and Marsh (2005). Though during the 2007-2009 crisis we have seen very negative basis for investment grade spreads.

¹⁵There are some technical differences between the CDX contract specification and a portfolio of CDS. Most notably, the fact that the CDX is settled upfront based on a fixed running coupon spread (see Collin-Dufresne (2008)). This has only a minor impact on our analysis below, however and therefore we ignore them.

where $R_i(t)$ is the recovery rate on firm i when it defaults at time t .

In addition to the CDX index, tranches based on the index have also been created. Each tranche is characterized by its attachment points. For the CDX index, the different tranches are 0-3% (the equity tranche), 3-7% (mezzanine), 7-10%, 10-15%, 15-30% (senior), and 30-100% (super-senior). The buyer of protection pays a periodic premium in return for compensation if there are losses on the CDX index that fall within the range of the particular tranche. Once again, the spread is determined by equating the present value of the protection leg and premium leg.

The tranche loss as a function of portfolio loss is

$$T_j(L(t)) \equiv T_{K_{j-1}, K_j}(L(t)) = \max[\min(L(t), K_j) - K_{j-1}, 0] \quad (11)$$

$$= \max[L(t) - K_{j-1}, 0] - \max[L(t) - K_j, 0]. \quad (12)$$

The initial value of the protection leg on tranche- j is

$$Prot_j(0, T) = E^Q \left[\int_0^T e^{-rt} dT_j(L(t)) \right] \quad (13)$$

In terms of the tranche spread S_j , the initial value of the premium leg on tranche- j (except for the equity and super-senior tranches) is

$$Prem_j(0, T) = S_j E^Q \left[\sum_{m=1}^M e^{-rt_m} \int_{t_{m-1}}^{t_m} du (K_j - K_{j-1} - T_j(L(u))) \right]. \quad (14)$$

There are two standard practices followed for the premium on the equity tranche. One is a so-called “full running premium” as in equation (14). In the other approach, which is more common, the equity tranche premium has a so-called “up-front premium” U combined with a set running premium of 500bps:

$$Prem_1(0, T) = UK_1 + 0.05 E^Q \left[\sum_{m=1}^M e^{-rt_m} \int_{t_{m-1}}^{t_m} du (K_1 - K_0 - T_1(L(u))) \right]. \quad (15)$$

The reason for these two different practices is that the pricing of the equity tranche using equation (14) is known to be very sensitive to the timing of defaults. To see this intuitively, note that the quarterly payment drops from $(0.03) \left(\frac{S}{4}\right)$ before any defaults to $(0.03 - \frac{(1-R)}{125}) \left(\frac{S}{4}\right)$ after the first default. For a standard recovery rate of $R = 0.4$, this is approximately a 16% drop in payments. The fact that the market is so concerned about the timing of defaults for the equity tranche that it has created a different cash flow for it emphasizes that the market is very aware of this sensitivity, and no doubt calibrating their models to incorporate information regarding this.

Finally, the super-senior tranche premium is specified by

$$Prem_6(0, T) = S_6 E^Q \left[\sum_{m=1}^M e^{-rt_m} \int_{t_{m-1}}^{t_m} du (K_6 - K_5 - n_u R - T_6(L(u))) \right]. \quad (16)$$

We note that high recoveries ‘hurt’ the super-senior tranches as the principal gets reduced by the amount recovered.¹⁶ Following CJS, we first focus on the period September 2004 to September 2007. We then investigate how the model performs during the crisis.

3 Data and calibration

3.1 Data

Our primary data include S&P 500 European option prices, the CDX North American Investment-Grade Index, and tranche spreads written on this index for maturities of one to five years. The option data are from OptionMetrics. The CDX index and tranche data are from J.P. Morgan. Our sample consists of two periods: the pre-crisis period (September 21, 2004 to September 20, 2007) and the crisis period (September 21, 2007 to September 20, 2008). The composition of the CDX index is refreshed every six months. The pre-crisis period includes data from Series 3 through Series 8, whereas the crisis period includes data from Series 9 through Series 10. To simplify the calibration and reduce bid-ask bounce error, we use only the closing quotes on every Wednesday in the sample. We remove options with maturities less than one year because the slope of the volatility smile of these options are highly sensitive to smaller but more frequent jumps which have little impact on the risk-neutral distribution of market values over long horizons.¹⁷ There are 206 weeks in our sample. The available options vary across moneynesses and maturities (between 1 and 3 years). While we use this option data to determine the time series of our state variables, we also use five year option data from CJS to calibrate the parameters of the model.

To identify asset betas, idiosyncratic volatilities and leverage ratios of firms in the collateral pool, we use daily CRSP and quarterly/annual Compustat data from 2000 to 2008. We obtain daily risk-free return from Ken French’s website. Treasury yields are from FRED, the database of the Federal Reserve Bank of St. Louis.

¹⁶Technically if recovery amounts exceeded the notional of the super-senior tranche then the notional of the adjacent senior tranche should be written down and so on. In practice, this is very unlikely to happen. The model is easily amended to account for that, however.

¹⁷That is, to jointly match short and long-maturity options would necessitate us adding additional jumps which are smaller in magnitude but more frequent. As has been shown in the previous literature, these jumps have negligible impact on the prices of long dated options.

3.2 Calibration of the option return model

First, we calibrate our SVDCJ option model to the 5-year index options (Series 4) from CJS. Using estimates from BCJ, we set the correlation coefficient between the market return and the market variance to $\rho = -0.48$ and the volatility parameter of the market variance to $\sigma_V = 0.2016$. Theory implies that these two parameters should remain the same under both the P-measure and Q-measures. We also choose the risk-free rate $r = 0.05$, and the average market dividend yield $\bar{\delta} = 0.04$. The rest of the parameters are chosen to minimize the percentage root mean square error (RMSE) of the 13 implied volatilities. We fix the initial values of the state variables to their long run means $V_0 = \bar{V}$ and $\delta_0 = \bar{\delta}$. The calibrated parameters correspond to “Estimate 1” in Table 1.

Parameter	Estimation 1	Estimation 2	Estimation 3	Estimation 4
ρ	-0.48	-0.48	-0.48	-0.48
σ_V	0.2016	0.2016	0.2016	0.2016
λ	0.1534	0.1608	0.1743	0.2465
ρ_q	0.0203	0.0199	-0.0509	-0.0576
μ_y	-0.2991	-0.2843	-0.4726	-0.3479
σ_y	0.2445	0.2441	0.4609	0.3915
\bar{V}	0.0037	0.0038	0.0132	0.0094
μ_V	0.0035	0.0033	0.0099	0.0056
κ_V	5.4368	5.3644	1.5442	2.1596
V_0	0.0037	0.0038	0.0132	0.0094
κ_δ	-0.5914	-0.5903	-0.4816	-0.4953
$\bar{\delta}$	0.04	0.04	0.04	0.04
σ_δ	0.0454	0.0423	0.0405	0.0304
ρ_1	-0.9054	-0.8968	-0.5056	-0.4135
ρ_2	-0.0032	-0.0036	-0.0078	-0.0066
μ_d	0.0002	0.0002	0.003	0.0007
σ_d	0.0007	0.0008	0.0006	0.0006
δ_0	0.04	0.04	0.04	0.04
r	0.05	0.05	0.05	0.05
y_C	0	-2	0	-2
λ_C^Q	0	0.00076	0	0.0066

Table 1: Calibration of market dynamics using two methods: The first method uses option data only, and sets the catastrophic jump size y_C and intensity λ_C^Q to zero. The second method uses both option data and the super-senior tranche spread, and calibrates the catastrophic jump intensity to minimize percentage RMSE (jump size set to $y_C = -2$). Both methods are applied to the pre-crisis (Estimation 1 and 2) and the crisis (Estimation 3 and 4) data.

A few parameter estimates are worth mentioning. Consistent with BCJ, we find an economically insignificant point estimate for ρ_q . More importantly, the point estimate for the mean reversion coefficient on dividend yield κ_δ is negative, implying non-stationarity. That is, consistent with the findings of Carr and Wu (2003), fitting long maturity option data requires

a violation of one of the assumptions underlying the central limit theorem. With that said, we show in the robustness section that our main results regarding the CDX tranche spreads are not driven by this result, and indeed similar results are obtained if we assume the dividend yield is constant.

Figure 2 shows the fit of the model for our calibration as well as the corresponding implied risk-neutral distribution. It is apparent that the model does a very good job at fitting the sample of 13 long-dated implied volatilities obtained by CJS. We emphasize that since the parameters are similar to those of BCJ, and that the stochastic dividend yield mainly impacts long-dated option prices, we expect the model to also fit short dated S&P 500 futures options well.

The second column denoted by “Estimate 2” in Table 1 corresponds to a second calibration where we set the size of the catastrophic jump to $y_C = -2.0$ (which corresponds to a -87% drop in price) and then choose its intensity λ_C^Q to match the average spread on the super-senior tranche for pre-crisis period. The recovery rate for those firms that default due to this jump is specified as 20% to capture the empirical finding that recovery rate is procyclical. (See, for example, Altman et al (2005)). We will discuss this in greater detail below, but it is worth emphasizing, as shown in Figure (2), that both calibrations imply almost identical option prices for those strikes that are actually traded, that is, for option strikes with moneyness greater than 0.7. Indeed, calibrating market dynamics to match the super-senior tranche impacts state prices significantly only for moneyness levels around 0.2 or lower. This is our first indication that one cannot ‘extrapolate’ the information in option prices to deduce information regarding the super-senior tranche. Instead, we calibrate our model to match the super-senior tranche, and investigate the model implications for the other CDX tranches.

The third and fourth columns are parameter values for the model calibrated to match the 5-year options in September, 2008. The risk-neutral catastrophe jump intensity is calibrated to the average spread of the super-senior tranche of CDX Series 10 (March, 2008 to September, 2008). The corresponding implied volatility curve and risk-neutral dynamics are shown in Figure 3. Comparing these results to those of Figure 2, we find rather dramatic changes to market dynamics during the crisis.

3.3 Calibration of asset value process

The most crucial distinction between our calibration procedure and that of CJS is the following: CJS assumes that idiosyncratic risk is driven by a diffusion process, and then set the location of the default boundary at each date to perfectly match the five-year CDX index. In contrast,

we assume idiosyncratic risk is driven by both diffusions and jumps, and calibrate our model to match the entire term structure of CDX indices. In particular, as in CJS, we use equity data to specify the idiosyncratic diffusion parameter. However, we set the idiosyncratic jump size to $y_i = -2.0$ (a value which basically guarantees default), and then calibrate the one-year, two-year, three-year, and four-year intensities¹⁸ on the idiosyncratic jump, and the location of the default boundary, to perfectly match the CDX indices at all five of these horizons. We argue that including jumps is essential, since it is well-documented that diffusion-based structural models of default for investment grade firms fail badly at capturing spreads (i.e, risk-neutral expected losses) at short maturities. As we demonstrate below, the pricing of tranche spreads is very sensitive to matching these (risk-neutral) expected loss rates.

To estimate asset beta, we take a weighted average of equity beta and debt beta. For each publicly traded firm in the on-the-run CDX series, equity beta is estimated using backward-looking 5-year daily equity returns prior to the first trading date of the series. As shown in Table 2, equity betas varied considerably over time, ranging from 0.82 to 0.99, and increased monotonically as the crisis approached.¹⁹ Moreover, market volatility increased, and idiosyncratic volatility decreased over our sample period. The time series of the fraction of systematic versus idiosyncratic risk implied from our model calibrated to both options, CDX index and tranches are shown in Figure 5. As discussed previously, the distribution of losses becomes more peaked around expected losses as the fraction of idiosyncratic risk increases. This is shown in Figure 4.

We estimate the beta of debt by regressing the excess returns of LQD (an ETF of investment graded corporate bonds) on the excess returns of S&P 500 for the same window of each series as we estimate the equity beta. The leverage ratio is defined as the book debt divided by the sum of book debt plus market equity, where book debt is from quarterly Compustat, using the short-term debt (DATA45) plus the long-term debt (DATA51). If any of these two numbers are missing, we use the corresponding items in the annual data. The cross-sectional average of leverage ratios are also displayed in Table 2.

3.4 Time Series of default boundary

Early papers (e.g., Merton (1974)) specified the default boundary to equal the face value (or book value) of debt. However, there are at least three strands of literature that suggest the location of the default boundary may be significantly lower. First, the literature on optimal

¹⁸We set the five-year intensity equal to the four-year intensity. This way, we have five free parameters to match the five CDX index spreads.

¹⁹We emphasize that the change in beta over time is due in part to a change in composition of the index.

Series	Period	Equity Beta	Leverage Ratio	Market Volatility	Idiosyncratic Asset Volatility
3	9/2004-3/2005	0.82	0.36	10.34	27.08
4	3/2005-9/2005	0.83	0.36	10.38	25.29
5	9/2005-3/2006	0.87	0.33	10.02	23.86
6	3/2006-9/2006	0.92	0.33	11.35	21.84
7	9/2006-3/2007	0.94	0.32	9.80	20.93
8	3/2007-9/2007	0.94	0.32	15.67	19.90
9	9/2007-3/2008	0.98	0.31	21.86	18.64
10	3/2008-9/2008	0.99	0.29	23.42	18.61

Table 2: Estimates of equity beta, leverage ratio, market and idiosyncratic volatility for each six month period that a given series was on the run. The betas are estimated via CAPM regression, based on backward-looking five year daily stock returns and S&P 500 returns. When estimating equity betas and leverage ratio, we include all publicly traded firms in the collateral pool.

capital structure (e.g., Leland (1994)) predicts that equityholders may find it optimal to cover debt payments via equity issuance in order to keep their implicit option to the firm's cash flows alive even if firm value falls well below the face value of debt. Second, since historical recovery rates have averaged approximately 45%, default at the face value of debt would imply bankruptcy costs of 55%, a number which is difficult to believe, since one would expect workouts to occur more often if so much value was on the line. Furthermore, direct estimates of bankruptcy costs are on the order of 20% (e.g., Andrade and Kaplan (1998)). Combining these two numbers gives a prediction for the ratio of default boundary to face value of debt (F) equal to $\frac{A_B}{F} = \frac{0.45}{1-0.2} \approx 0.56$. Third, Leland (2004) finds that the standard structural model of default under the historical measure is generally consistent with historical default rates (for maturities greater than three years) if the default boundary is approximately 70% of the face value of debt. Similar results are obtained by Davydenko (2008), who estimates this ratio to be in the range $\frac{A_B}{F} \in (56\%, 70\%)$.

As shown in Figure 6, the ratio of our calibrated default boundary to book value of debt varies significantly over our sample period, monotonically increasing from 57% at the beginning of our sample to nearly 100% during the crisis.

In the calibrations below, we set the (risk-neutral expected) recovery rate to 40% for individual defaults, consistent with market convention, and slightly below the 45% historical recovery rate on investment grade senior unsecured debt. However, there is evidence that recovery is procyclical (e.g., Altman et al (2005)). As a recent example, the recovery rate on

GM debt, whose default occurred during the crisis, was approximately 10%. Therefore, in our benchmark model we set recovery rate to 20% in the event of a market collapse.

Finally, we note that the average risk-neutral intensity associated with market collapse is $\lambda_C^Q = 0.00076$. That is, the market anticipates a catastrophic event to occur less than once per thousand years. Although this intensity is miniscule, it has a significant impact on the size of the more senior tranche spreads.

4 Results

4.1 Average tranche spreads: pre-crisis period

In this section we report six sets of tranche spreads averaged over the pre-crisis period September 2004 - September 2007:

- The historical values;
- Benchmark model: Catastrophic jumps calibrated to match the super-senior tranche spread; Idiosyncratic jumps and default boundary calibrated to match the 1-year, 2-year, 3-year, 4-year and 5-year CDX index spreads.
- $\lambda_C^Q = 0$: No catastrophic jumps; Idiosyncratic jumps and default boundary calibrated to match the 1-year, 2-year, 3-year, 4-year and 5-year CDX index spreads;
- $\lambda_i^Q = 0$: Catastrophic jumps calibrated to match the super-senior tranche spread; No idiosyncratic jumps; Default boundary calibrated to match only the 5Y CDX index.
- $\lambda_C^Q = 0, \lambda_i^Q = 0$: No catastrophic jumps; No idiosyncratic jumps; Default boundary calibrated to match only the 5Y CDX index;
- The results reported by CJS

The results are given in Table 3. As shown in the last row, our framework generates errors that are an order of magnitude smaller than errors reported by CJS. Even in absolute terms, our framework performs well. As stated earlier, given the results of the previous literature (e.g., Longstaff and Rajan (2008)), we are confident that the fit could be improved even further if we were to calibrate market dynamics to match, say, mezzanine tranche spreads, rather than pricing them out-of-sample as we do here.

Here we use the results of the other three models to provide some intuition for why our model performs so much better than that of CJS. We begin with the case closest to CJS;

	0-3%	3-7%	7-10%	10-15%	15-30%	30-100%	0-3% Upfirt
data	1472	135	37	17	8	4	0.34
benchmark	1449	113	25	13	8	4	0.33
$\lambda_C^Q = 0$	1669	133	21	6	1	0	0.40
$\lambda_i^Q = 0$	1077	206	70	32	12	4	0.22
$\lambda_C^Q = 0, \lambda_i^Q = 0$	1184	238	79	31	6	0	0.26
CJS	914	267	150	87	28	1	na
$\frac{ CJS-Data }{ Benchmark-Data }$	24.3	6	9.4	17.5	∞	∞	

Table 3: Historical and model estimated average tranche spreads over the time interval September 2004 - September 2007 for four different models: i) benchmark, ii) benchmark without catastrophic jump, iii) benchmark without idiosyncratic jumps, iv) benchmark without either catastrophic jump or idiosyncratic jumps. For comparison, we also report the results of CJS.

the model with no catastrophic and no idiosyncratic jumps ($\lambda_C^Q = 0, \lambda_i^Q = 0$). As with CJS, this model is calibrated to match option prices and the 5 year CDX index (and thus, it matches 5-year risk neutral expected losses). However, without idiosyncratic jumps, this model generates short horizon CDX index spreads that are well below observation. As such, expected losses are “backloaded”, implying that the buyer of equity protection pays too much premium for too long, in turn biasing down the estimate for the equity tranche spread. Since the model is calibrated to match expected losses, this downward bias on the equity tranche spread automatically biases upward spreads on the more senior tranches.

To see how important this backloading issue is, we determine the CDX index spreads for maturities of 1-4 years for each of the models. The results are displayed in Table 4. For example, we find that the ($\lambda_C^Q = 0, \lambda_i^Q = 0$) model predicts credit spreads of 0bp and 3bp at maturities of 1 and 2 years, respectively, well below the market quotes of 13bp and 20bp.

Adding idiosyncratic jumps calibrated to match short horizon credit spreads not only solves the backloading problem, but it also generates a five-year loss distribution that is more peaked about the risk-neutral expected losses of 2.4%. Indeed, the standard deviation of the loss distribution without idiosyncratic jumps ($\lambda_C^Q = 0, \lambda_i^Q = 0$) has a standard deviation of 2.9%, whereas the standard deviation of the loss distribution with idiosyncratic jumps ($\lambda_C^Q = 0$) has a standard deviation of 1.7%. The implication is that adding idiosyncratic jumps generates a smaller probability of losses falling into the more senior tranches, pushing their spreads down, and in turn the equity tranche spread up.

Thus, the downward bias of the equity tranche spread in the ($\lambda_C^Q = 0, \lambda_i^Q = 0$) model has two sources: one due to backloading, and one due to an error in the idiosyncratic risk/systematic risk composition. In an attempt to decompose this bias into its components, we approximate the backloading bias by treating losses as deterministic, and equal to the risk-neutral expected

losses implied in Table 4. Note that this approach will produce an equity tranche spread that is biased upward for each model, since this approach does not cut off losses at the equity detachment point of 3%. As such, we focus not on the *levels* generated by this approximation, but rather on the *difference in levels* across these two models. Specifically, we assume that the term structure of loss rates is piecewise constant over the intervals (0-1 year), (1-2 year), (2-3 year), (3-4 year), and (4-5 year). The loss rates are chosen to perfectly match the implied term structure of credit spreads in Table 4. After these loss rates are determined, they are used to calculate the protection leg and premium leg of the tranche using equations (13) and (14). Calibrated to actual data on the term structure of CDX spreads (row 1 in table 4), we find the equity tranche spread equals 1929bp. In contrast, calibrated to the term structure of CDX spreads implied by our $(\lambda_C^Q = 0, \lambda_i^Q = 0)$ model (row 5 in table 4), the equity tranche spread is only 1692bp. While both of these levels are biased upward by our approximation technique, the difference between these two estimates, 237bp, provides an estimate for the amount of downward bias in the equity tranche spread that can be attributed to backloading.

If we start with the case $(\lambda_C^Q = 0, \lambda_i^Q = 0)$, and then add idiosyncratic jumps to match the term structure of CDX index spreads, we get the model $(\lambda_C^Q = 0)$, the results of which are shown in the third line of Table 3. Interestingly, we find in this case that the predicted equity tranche spread is actually too high – implying a problem opposite to that reported by CJS. Since the backloading problem has been resolved in this case, these results imply that the $(\lambda_C^Q = 0)$ model is too peaked – that is, has too high a ratio of idiosyncratic risk to systematic risk. But this problem is easily resolved by adding a catastrophic jump calibrated to match the super-senior tranche spread (which produces our Benchmark model). As shown previously, including a catastrophic jump has virtually no impact on the model’s ability to match option prices.

In summary then, we find that in order to estimate tranche spreads, it is necessary that the model be calibrated to match the term structure of credit spreads. Specifying a model with idiosyncratic dynamics driven only by diffusive risks generates a model where the timing of defaults is backloaded. This causes counter-factually low spreads/losses at short maturities, which in turn biases down the equity tranche spread. In addition, the super-senior tranche spread (and therefore, spreads on other senior tranches) cannot be *extrapolated* from option prices alone. However, spreads on other tranches can be *interpolated* reasonably well given option prices and super-senior tranche spreads. We conclude, in contrast with the results of CJS, that S&P 500 options and CDX tranche prices market can be fairly well reconciled within our arbitrage-free model. In that sense these two markets appear to be well integrated.

	1 year	2 year	3 year	4 year	5 year
Data	13	20	28	36	45
Benchmark	13	20	28	36	45
$\lambda_C^Q = 0$	13	20	28	36	45
$\lambda_i^Q = 0$	6	7	16	29	45
$(\lambda_C^Q = 0, \lambda_i^Q = 0)$	0	3	13	28	45

Table 4: Historical and model-estimated average CDX index spreads September 2004 - September 2007 for four different models: i) benchmark, ii) benchmark without catastrophic jump, iii) benchmark without idiosyncratic jumps, iv) benchmark without either catastrophic jump or idiosyncratic jumps.

4.2 Time-series of tranche spreads

In the previous section we showed that our benchmark model is able to fit average historical spreads across all tranches very well, once it is calibrated to match the term structure of CDX index spreads and the super-senior claim. In this section we investigate its time-series performance.

For Series 3 to Series 9, we keep the parameters for the option pricing model fixed as given in Table 1 “Estimates 2”. For Series 10, we use the parameter values from the column labeled “Estimate 4”. Then, for each week, we use option prices to identify the state variables V_t and δ_t . The time-series of the RMSE are provided in Figure 7. In addition, each week we calibrate the intensity of the catastrophic jump to match the super-senior tranche, and the idiosyncratic jump intensities to match the term structure of CDX index spreads with maturities of one-year to five-years. The time series for the CDX indices and the super-senior tranche used by our calibration exercise are shown in Figure 8.

With this calibration, we then estimate tranche spreads. The results (which we consider to be the main contribution of this paper) are given in Figure 9. Consistent with the time-averaged results, the picture reveals a dramatic improvement in fit relative to that proposed by CJS.²⁰ Note that the performance is somewhat better during the pre-crisis period than the crisis period. For example, the model-implied 7-10% mezzanine tranche spread is too low, and the 15-30% senior tranche spread is too high during the crisis. Interestingly, even if we set the catastrophic jump to zero, the implied senior tranche spread remains too high during the crisis period. This is due to the very high implied volatilities on the deep out of the money puts on the S&P 500 options.²¹ Once again, we are confident that we could improve the fit of these

²⁰We do not plot the super-senior tranche since it is fit perfectly by construction.

²¹One possible cause is that we specify market *equity* dynamics as log-normal to simplify equity option pricing, even though theoretically it would be better to model market *asset* dynamics as log-normal (Toft and Prucyk (1997)). While straightforward to model, such a framework would be computationally less tractable. We leave

tranche spreads if we were to use them as part of the calibration procedure, but that is not the purpose of this exercise. We feel that the overall fit demonstrated in Figure 9 is quite good, and very much consistent with the belief that the S&P 500 option market and CDX tranche market are well-integrated.

Of course, our results only show that at every point in time it is possible to find an arbitrage-free model that consistently prices (in the cross-section) S&P500 options, the CDX index and CDX tranches. However, our model is not time-consistent, as we recalibrate certain parameters (such as the crash jump intensities) which within the model are assumed to be constant. It would be interesting to see if the model can be made time consistent by, for example, allowing for time-varying intensities or partial learning as in Cogley and Sargent (2008). We leave these interesting questions for future research.

5 Robustness

For parsimony, many simplifying assumptions were made in the previous sections, such as i) firm homogeneity, ii) no changes in capital structure, iii) uncorrelated idiosyncratic shocks (i.e., “no industry effects”), iv) constant firm-level asset dividend yield, even though market equity dividend yield is stochastic, and v) constant interest rates. Here we show that our results are robust to these assumptions.

We investigate deviations from our benchmark model along several dimensions. For each deviation, we continue to calibrate the model to 5-year option implied volatilities, 1-5 year CDX indices, and the super-senior tranche spreads. Then we report model implied tranche spreads. The results are given in Table 5. Examples include:

- **Dynamic capital structure:** We assume that if a firm performs well, it will issue additional debt, in turn raising the default boundary (e.g., Goldstein, Leland and Ju (2001)). Specifically, we specify default boundary dynamics $A_B(t)$ via: $A_B(t + dt) = \max[A_B(t), cA(t)]$, where we set the constant c to the initial leverage ratio divided by 1.1. Thus, the default boundary remains constant until firm value increases by at least 10% from its initial value. The result is reported in the first row as ‘Dynamic capital structure’.
- **Stochastic asset dividend yield at the firm level:** We specify the firm payout ratio as $\delta_A(t) = \bar{\delta}_A + \xi(\delta_t - \bar{\delta})$, where $\bar{\delta}_A = 0.05$ is the average payout ratio, and $\xi = 0.7$ measures the correlation of dynamics of the firm payout ratio and market dividend-price ratio. We choose a value of ξ to be less than one to account for the fact that we anticipate equity

this interesting question for future research.

dividend yields to be more volatile than yields on the asset, since coupon payments are stickier than dividend payments.

- Constant market equity dividend yield: we specify market dynamics using the SVCJ option model so that both the market dividend price ratio and the firm payout ratio are constants in this scenario. The purpose of this robustness check is to emphasize that our results are not being driven by the (non-stationary) dynamics on the dividend yield.
- Heterogeneity in initial credit spreads: We use our model to back out the default boundaries for each firm based on their average 5-year CDS spreads in the on-the-run period of Series 4. The 5-year CDS spreads are from Datastream. The cross-sectional mean and the standard deviation of the log default boundaries are -1.59 and 0.344. Instead of specifying a homogenous initial leverage ratio for all firms as in the benchmark model, here we specify a distribution for the log default boundaries of the 125 firms using a normal distribution with the above parameters.
- Stochastic interest rates: We specify the spot rate to follow Vasiec (1977) interest rates and calibrate the model to the average term structure of treasury rates for maturities 3-month, 6-month, 1-year, 2-year, 3-year and 5-year for pre-crisis sample period. We assume the interest rate process is independent of any other random shock in the model to simplify the calibration. The option model with the Vasiec-type stochastic interest rate is still affine, so that we can use the FFT to solve for option prices.
- Industry Correlations: Our benchmark model assumes a CAPM-like structure, where there is only market and idiosyncratic risk. It is straightforward to include other sources of risks that are shared by only a fraction of the 125 firms, capturing the notion that industry correlations may be stronger than a CAPM calibration would predict. As a simple way to capture this feature, we assume that there are approximately two firms per industry with dynamics that are perfectly correlated. As such, instead of modeling 125 firms, we consider only 60 “industries”.

We see that, while these different assumptions do impact tranche spreads, they are of second-order importance compared to the need to calibrate the model to the term structure of CDX index spreads and the super-senior tranche spread. We interpret these findings to imply that our results are robust along many different dimensions.

	0-3%	3-7%	7-10%	10-15%	15-30%	30-100%	0-3% Upfirt
data	1472	135	37	17	8	4	0.34
benchmark	1449	113	25	13	8	4	0.33
Dynamic capital structure	1452	116	27	14	8	4	0.34
Stochastic firm payout	1441	122	29	14	9	4	0.33
SVCJ	1330	138	47	26	12	4	0.30
Heterogeneous initial credit spreads	1406	133	28	13	8	4	0.32
Stochastic short-term rate	1484	114	22	11	8	4	0.36
Industry Correlations	1370	153	31	16	10	5	0.31

Table 5: Robustness check

6 Conclusion

We examine the relative pricing of long-maturity S&P 500 option prices and CDX tranche spreads. We demonstrate the importance of calibrating the model to match the entire term structure of CDX index spreads because it contains pertinent information regarding both the timing of expected defaults and the specification of idiosyncratic dynamics. In particular, consistent with the previous literature (e.g., Jones, Mason and Rosenfeld (1984)), jumps must be added to idiosyncratic dynamics in order to explain credit spreads at short maturities. With this calibration approach, we find our model matches historical tranche prices extremely well, both in time series and in the cross section. In contrast to the conclusions of Coval, Jurek and Stafford (2009), we conclude that S&P 500 options and CDX tranche prices market can be reconciled within an arbitrage-free framework. In that sense these two markets appear to be well integrated.

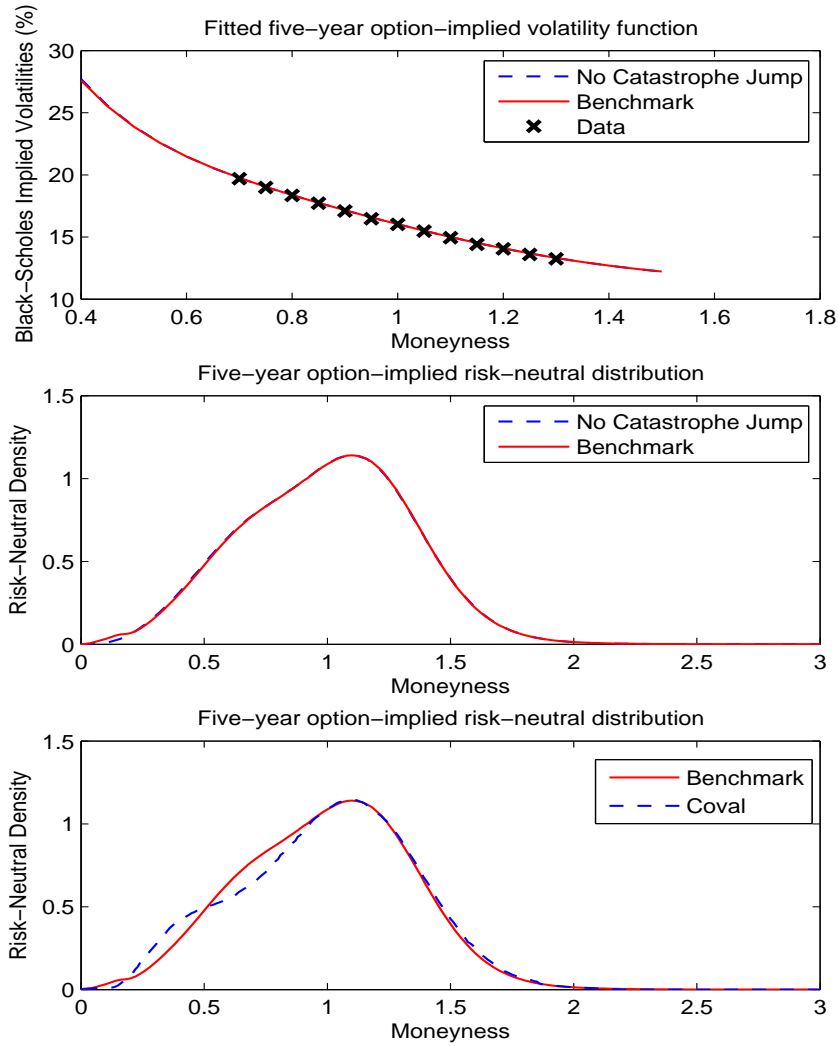


Figure 2: Market dynamics given in equations (1)-(3) are calibrated to match five year option prices obtained from CJS during the pre-crisis period. The parameters are specified in the columns labeled “Estimation 1” and “Estimation 2” of Table 1. Figure 2A shows the model-implied five year volatility surface and the actual option prices as a function of moneyness for both the benchmark case where there is a catastrophic jump, and the case where there is not. Figure 2B converts these two implied volatility surfaces into risk-neutral distributions for the five-year level of market value. For comparison, Figure 2C demonstrates the risk-neutral distribution from CJS.

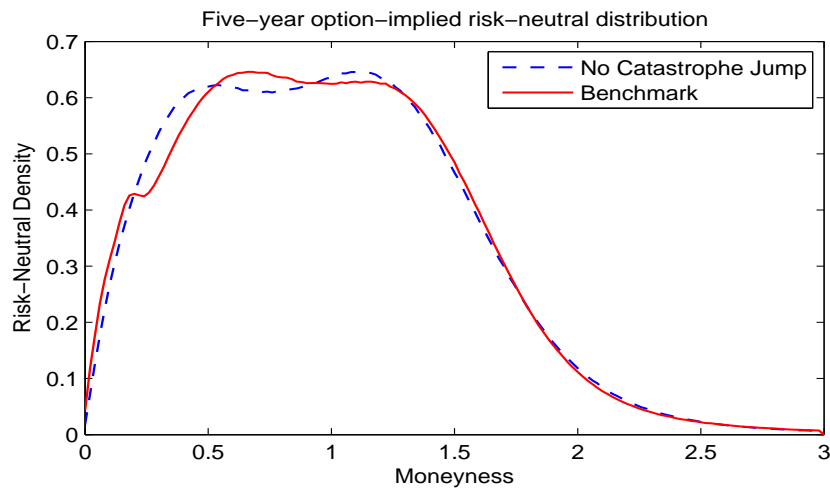
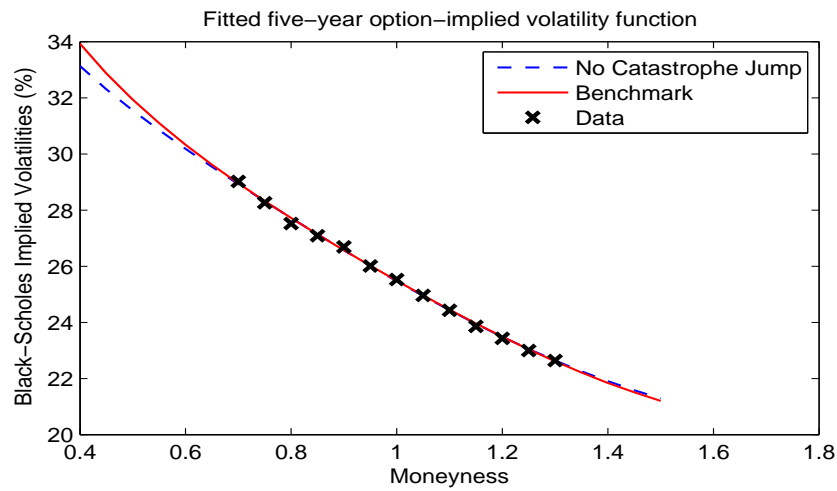


Figure 3: Market dynamics given in equations (1)-(3) are calibrated to match five year option prices obtained from CJS during the crisis period. The parameters are specified in the columns labeled “Estimation 3” and “Estimation 4” of Table 1. Figure 3A shows the model-implied five year volatility surface and the actual option prices as a function of moneyness for both the benchmark case where there is a catastrophic jump, and the case where there is not. Figure 3B converts these two implied volatility surfaces into risk-neutral distributions for the five-year level of market value.

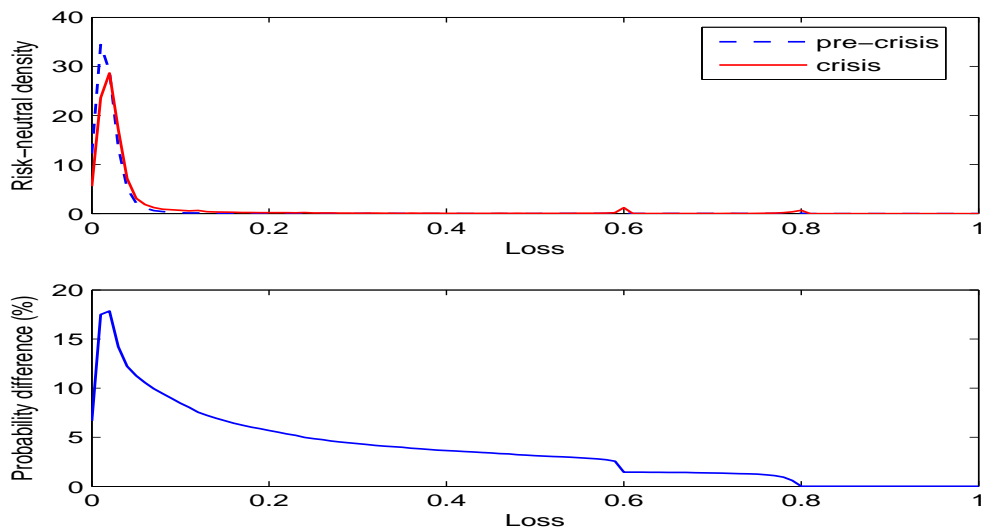


Figure 4: In Figure 4A, we plot the risk-neutral loss density for the pre-crisis and crisis periods. The crisis period has higher expected losses and a less-peaked distribution due to a larger proportion of risk being systematic. In Figure 4B, we plot the difference in the cumulative loss distributions for the crisis and pre-crisis periods.

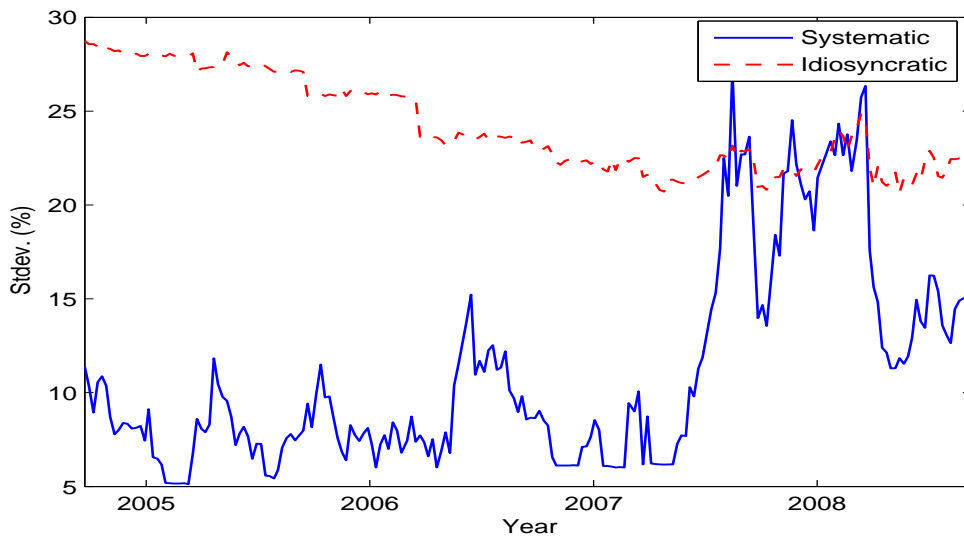


Figure 5: The time-series of the standard deviation of the systematic and idiosyncratic risk for the firm value. The systematic risk is measured as the asset beta times the standard deviation of the market diffusion and jumps implied by S&P 500 options and CDX super-senior tranche. The idiosyncratic risk is measured as the standard deviation of the idiosyncratic diffusion and jump implied by 1, 2, 3, 4, and 5 year CDX indices.

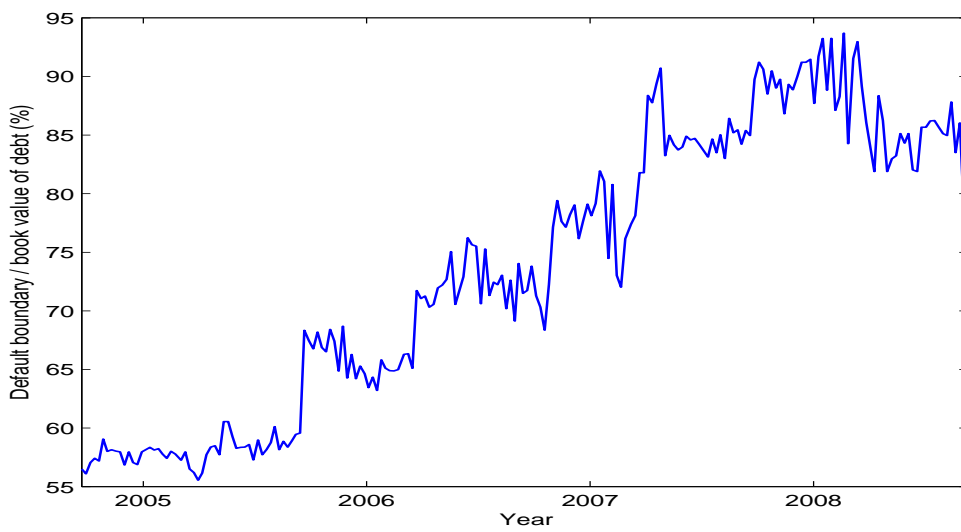


Figure 6: The time series of the ratio of implied default boundary to book value of debt. Many strands of literature predict that this ratio should be significantly less than unity.

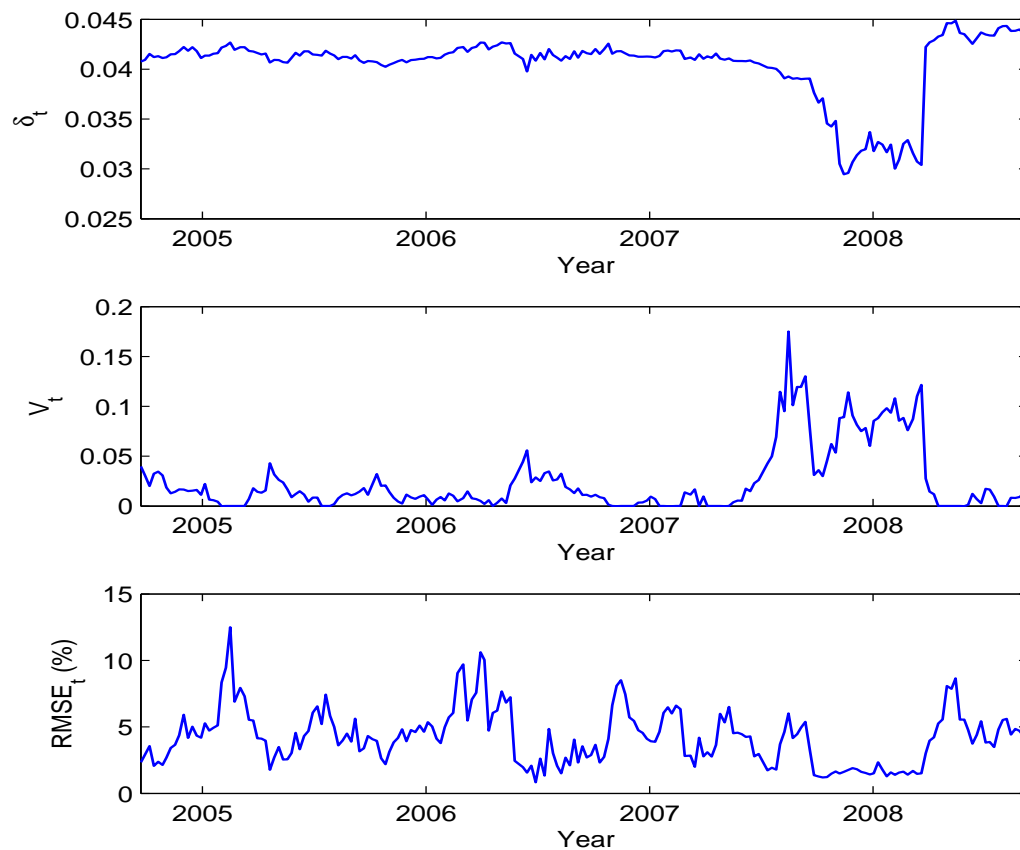


Figure 7: The time-series of the implied dividend yield δ_t , the implied market variance V_t , and the relative RMSE of the option implied volatilities. We fix the parameters that are chosen to match the 5-year options, and calibrate the two state variables, the dividend yield and the market variance for each date to minimizing the relative RMSE of the option with moneyness from 0.7 to 1 with 0.05 increments.

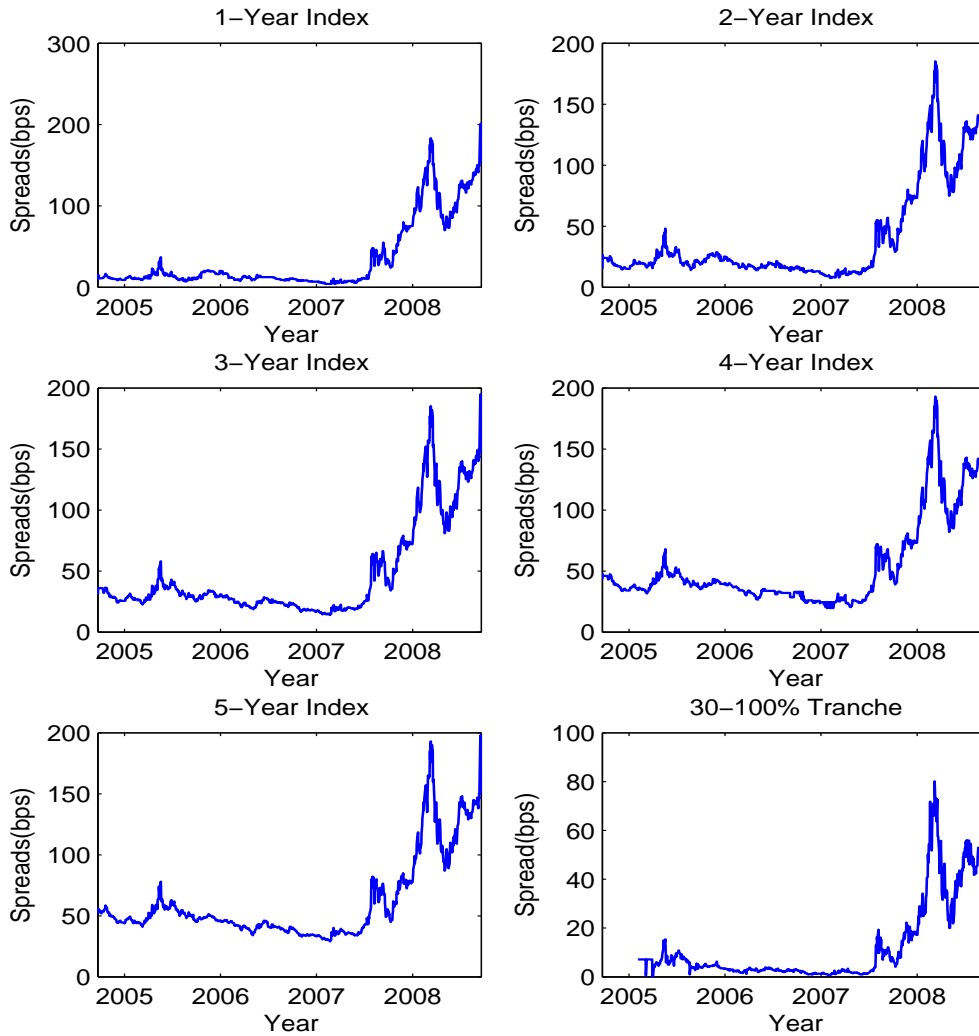


Figure 8: Historical time series of spreads for the one-year, two-year, three-year, four-year and five-year CDX indices, and the time series of the 30-100% super-senior tranche. Our benchmark model is calibrated to perfectly match these time series.

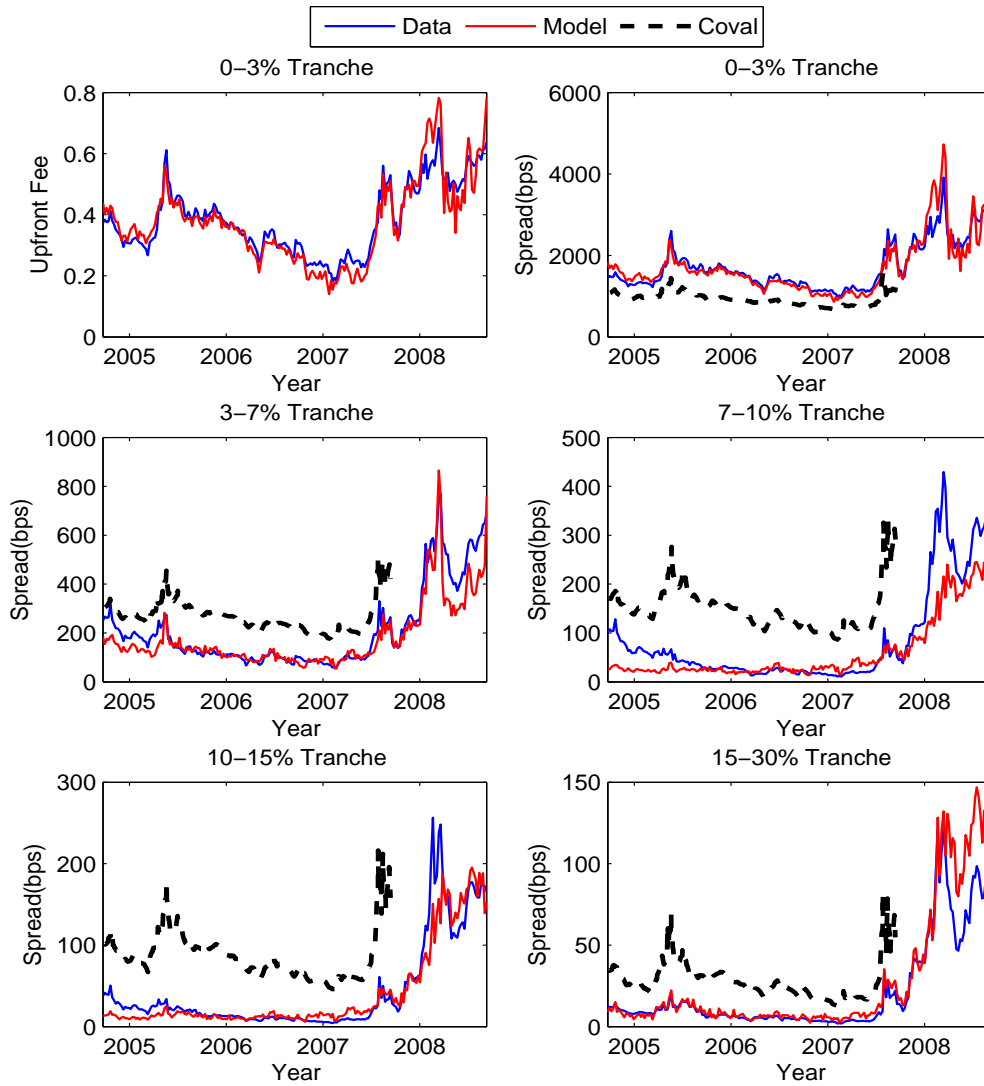


Figure 9: Three time series of spreads for the 0-3% up-front premium, 0-3% running premium, 3-7%, 7-10%, 10-15%, 15-30% CDX tranches. The smooth (blue) lines are the historical data. The dashed (red) lines are the spreads implied by our model. The dot-dashed (black) lines are the spreads implied by CJS.

References

- [1] E. I. Altman, B. Brady, A. Resti, and A. Sironi. The link between default and recovery rates: Theory, empirical evidence, and implications. *Journal of Business*, 78:2203–2227, 2005.
- [2] L. Andersen and J. Andreasen. Jump-diffusion processes: Volatility smile fitting and numerical methods for option pricing. *Review of Derivatives Research*, 4:1380–6645, 2000.
- [3] L. Andersen and J. Sidenius. Extensions to the gaussian copula: random recovery and random factor loadings. *Journal of Credit Risk*, 1(1), 2004.
- [4] D. Backus, S. Foresi, and L. Wu. Accounting for biases in black-scholes. *working paper*, 2004.
- [5] G. Bakshi, C. Cao, and Z. Chen. Empirical performance of alternative option pricing models. *Journal of Finance*, 52:2003–2049, 1997.
- [6] G. Bakshi, D. Madan, and F. X. Zhang. Investigating the role of systematic and firm-specific factors in default risk: Lessons from empirically evaluating credit risk models. *Journal of Business*, 79:1955–1988, 2006.
- [7] F. Black and J. C. Cox. Valuing corporate securities: some effects of bond indentures provisions. *Journal of Finance*, 31:351–367, 1976.
- [8] F. Black and M. Scholes. The pricing of options and corporate liabilities. *Journal of Political Economy*, 81:637–654, 1973.
- [9] R. Blanco, S. Brennan, and I. W. Marsh. An empirical analysis of the dynamic relation between investment-grade bonds and credit default swaps. *Journal of Finance*, 60:2255–2281, 2005.
- [10] D. T. Breeden and L. Litzenberger. Prices of stochastic contingent claims implicit in option prices. *Journal of Business*, 51:621–651, 1978.
- [11] M. Broadie, M. Chernov, and M. Johannes. The impact of jumps in volatility and returns. *Journal of Finance*, 62:1453–1490, 2007.
- [12] J. Campbell and J. Cochrane. By force of habit: A consumption-based explanation of aggregate stock return behavior. *Journal of Political Economy*, 107:205–251, 1999.

- [13] P. Carr and D. Madan. Option valuation using the fast fourier transform. *Journal of Computational Finance*, 2:61–73, 1998.
- [14] P. Carr and L. Wu. The finite moment log stable process and option pricing. *Journal of Finance*, 58:753–777, 2003.
- [15] L. Chen, P. Collin-Dufresne, and R. Goldstein. On the relation between the credit spread puzzle and the equity premium puzzle. *Review of Financial Studies*, 22:3367–3409, 2009.
- [16] T. Cogley and T. Sargent. Anticipated utility and rational expectations as approximations of bayesian decision making. *International Economic Review*, 49:1:185–221, 2008.
- [17] P. Collin-Dufresne, R. S. Goldstein, and J. Helwege. Is credit event risk priced? modeling contagion via the updating of beliefs. *working paper*, 2003.
- [18] J. D. Coval, J. W. Jurek, and E. Stafford. Economic catastrophe bonds. *American Economic Review*, 99(3):628–666, 2009.
- [19] K. M. Cremers, J. Driessen, and P. Maenhout. Explaining the level of credit spreads: Option-implied jump risk premia in a firm value model. *The Review of Financial Studies*, 21:2210–2242, 2008.
- [20] S. R. Das, D. Duffie, N. Kapadia, and L. Saita. Common failings: How corporate defaults are correlated. *Journal of Finance*, pages 93–117, 2007.
- [21] S. R. Das, L. Freed, G. Geng, and N. Kapadia. Correlated default risk. *Journal of Fixed Income*, pages 7–32, 2006.
- [22] S. R. Das and R. Sundaram. Of smiles and smirks: A term structure perspective. *Journal of Financial and Quantitative Analysis*, 34:211–240, 1999.
- [23] S. A. Davydenko. When do firms default? a study of the default boundary. *working paper*, 2007.
- [24] D. Duffie and N. Garleanu. Risk and valuation of collateralized debt obligations. *Financial Analysts Journal*, 57n1:41–59, 2001.
- [25] D. Duffie, J. Pan, and K. Singleton. Transform analysis and option pricing for affine jump-diffusions. *Econometrica*, 68:1343–1376, 2000.

- [26] D. Duffie, L. Saita, and K. Wang. Multi-period corporate default prediction with stochastic covariates. *Journal of Financial Economics*, 83:635–665, 2007.
- [27] B. Dumas, J. Fleming, and R. Whaley. Implied volatility functions: Empirical tests. *Journal of Finance*, 53:2059–2106, 1998.
- [28] Y. Eom, J. Helwege, and J.-Z. Huang. Structural models of corporate bond pricing. *RFS*, 17:499–544, 2004.
- [29] B. Eraker. Do stock prices and volatility jump? reconciling evidence from spot and option prices. *Journal of Finance*, 59:1367–1403, 2004.
- [30] B. Eraker, M. S. Johannes, and N. Polson. The impact of jumps in volatility and returns. *Journal of Finance*, 58:1269–1300, 2003.
- [31] K. Giesecke and L. R. Goldberg. A top down approach to multi-name credit. *working paper*, 2005.
- [32] R. Goldstein, N. Ju, and H. Leland. An ebit-based model of dynamic capital structure. Forthcoming, *Journal of Business*, 2001.
- [33] J. Hull and A. White. Valuation of a cdo and an nth to default cds without monte carlo simulation. *Journal of Derivatives*, 12n4:8–23, 2004.
- [34] P. E. Jones, S. P. Mason, and E. Rosenfeld. Contingent claims analysis of corporate capital structures: An empirical investigation. *Journal of Finance*, v39 n3 July:611–25, 1984.
- [35] P. Jorion and G. Zhang. Good and bad credit contagion: Evidence from credit default swaps. *Journal of Financial Economics*, 84:860–883, 2007.
- [36] H. Leland. Predictions of default probabilities in structural models of debt. *Journal of Investment Management*, 2(2), 2004.
- [37] D. X. Li. On default correlation: A copula function approach. *Journal of Fixed Income*, v9n4:43–54, 2000.
- [38] F. A. Longstaff, S. Mithal, and E. Neis. Corporate yield spreads: Default risk or liquidity? new evidence from the credit default swap market. *Journal of Finance*, 60:2213–2253, 2005.

- [39] F. A. Longstaff and A. Rajan. An empirical analysis of the pricing of collateralized debt obligations. *working paper*, 2006.
- [40] R. C. Merton. On the pricing of corporate debt: The risk structure of interest rates. *Journal of Finance*, 29:449–470, 1974.
- [41] A. Mortensen. Semi-analytical valuation of basket credit derivatives in intensity-based models. *Journal of Derivatives*, v13n4:8–26, 2006.
- [42] J. Pan. The jump-risk premia implicit in options: evidence from an integrated time-series study. *Journal of Financial Economics*, 63:3–50, 2002.
- [43] M. Rubinstein. Implied binomial trees. *Journal of Finance*, 49:771–818, 1994.
- [44] K. Toft and B. Prucyk. Options on leveraged equity: Theory and empirical tests. *Journal of Finance*, 52:1151–1180, 1997.
- [45] O. Vasicek. Probability of loss on a loan portfolio. *Moody's-KMV Research document*, 1987.
- [46] C. Zhou. An analysis of default correlation and multiple defaults. *The Review of Financial Studies*, v14(2):555–576, 2001.

7 Appendices

7.1 Fundamental of the Mean-Diverting Dividend Yield

Slightly modifying their notation, Campbell and Cochrane (1999) specify the utility function of the representative agent in an exchange economy as

$$U(C_t, \widehat{C}_t, t) = e^{-\alpha t} \frac{(C - \widehat{C})^{1-\gamma} - 1}{1-\gamma}, \quad (17)$$

where \widehat{C} is an exogenous habit. Campbell-Cochrane define the surplus consumption ratio as $S \equiv \left(\frac{C-\widehat{C}}{C}\right)$. For convenience, they also define the logarithms of consumption and surplus consumption via

$$c \equiv \log C \quad \text{and} \quad s \equiv \log S.$$

Because the dividend is perishable and there are no investment opportunities, it follows that in equilibrium consumption equals the dividend payment. Further, the pricing kernel is equal to the marginal utility of the representative agent:

$$\begin{aligned} \Lambda_t &= U_C(C_t, \widehat{C}_t, t) \\ &= e^{-\alpha t} (C - \widehat{C})^{-\gamma} \\ &= e^{-\alpha t} e^{-\gamma s} e^{-\gamma c}. \end{aligned} \quad (18)$$

In continuous time their log-consumption and surplus consumption processes are:

$$dc = g_c dt + \sigma_c dz \quad (19)$$

$$ds = \kappa(\bar{s} - s) dt + \sigma \left[\frac{1}{\bar{S}} \sqrt{1 - 2(s - \bar{s})} - 1 \right] dz, \quad (20)$$

where the constants \bar{S} and s_{max} are

$$\bar{S} \equiv \sigma \sqrt{\frac{\gamma}{\kappa}} \quad (21)$$

$$s_{max} \equiv \bar{s} + \frac{1}{2} (1 - \bar{S}^2). \quad (22)$$

From Ito's lemma, pricing kernel dynamics follow

$$\frac{d\Lambda}{\Lambda} = -r dt - \frac{\gamma\sigma}{\bar{S}} \sqrt{1 - 2(s - \bar{s})} dz. \quad (23)$$

Girsanov's theorem then implies that one can identify the risk-neutral measure by rewriting dynamics in terms of

$$dz^Q = dz + \left(\frac{\gamma\sigma}{\bar{S}} \sqrt{1 - 2(s - \bar{s})} \right) dt \quad (24)$$

Plugging this into equation (20), we find that the coefficient in the drift multiplying the term linear in s goes from $-\kappa$ under the historical measure to $+\kappa$ under the risk-neutral measure. This non-stationary feature motivates our market dynamics in equations (1)-(3).

7.2 Numerical solution of the long-term option model

We add two ingredients to the standard SVCJ option model. First, stochastic dividends are introduced to match the long-term options. Second, a catastrophe jump is introduced to match the super-senior tranche (30-100%) of the CDX. The dynamics are

$$\begin{aligned} dm_t &= \left(r - \delta_t - \bar{\mu}_y \lambda^Q - (e^{y_C} - 1) - \frac{1}{2} V_t \right) dt + \sqrt{V_t} dw_1^Q + y dq + y_C dq_C \\ dV_t &= \kappa_V (\bar{V} - V_t) dt + \sigma_V \sqrt{V_t} (\rho dw_1^Q + \sqrt{1 - \rho^2} dw_2) + y_V dq \\ d\delta_t &= \kappa_\delta (\bar{\delta} - \delta_t) dt + \sigma_\delta \sqrt{V_t} (\rho_1 dw_1^Q + \rho_2 dw_2 + \sqrt{1 - \rho_1^2 - \rho_2^2} dw_3) + y_\delta dq. \end{aligned}$$

The stochastic dividend yield δ_t is allowed to be negative, meaning financing. Since the dynamics are affine, the moment generating function of the log market is

$$\phi_T(u) = \mathbb{E}^Q [e^{um_T}] = e^{A(T) + um_0 + B(T)V_0 + C(T)\delta_0},$$

where $A(t), B(t), C(t)$ solves the ODEs,

$$\begin{aligned} A'(t) &= (r - \bar{\mu}_y \lambda^Q - (e^{y_C} - 1) \lambda_C^Q) u + \kappa_V \bar{V} B(t) + \kappa_\delta \bar{\delta} C(t) + \lambda^Q \mathbb{E} \left[e^{uy + B(t)y_V + C(t)y_\delta} - 1 \right] \\ &\quad + (e^{uy_C} - 1) \lambda_C^Q \\ B'(t) &= -\frac{1}{2} u + \frac{1}{2} u^2 - \kappa_V B(t) + \frac{1}{2} \sigma_V^2 B(t)^2 + \frac{1}{2} \sigma_\delta^2 C(t)^2 + \rho \sigma_V B(t) u \\ &\quad + (\rho \rho_1 + \rho_2 \sqrt{1 - \rho^2}) \sigma_V \sigma_\delta B(t) C(t) + \rho_1 \sigma_\delta u C(t) \\ C'(t) &= -u - \kappa_\delta C(t), \end{aligned}$$

where

$$\mathbb{E} \left[e^{uy + B(t)y_V + C(t)y_\delta} - 1 \right] = \frac{e^{u\mu_y + \frac{1}{2}u^2\sigma_y^2 + C(t)\mu_d + \frac{1}{2}C(t)^2\sigma_d^2}}{1 - \mu_V(B(t) + u\rho_q)} - 1.$$

We can apply the FFT to the characteristic function $\phi_T(iv)$ to get the risk-neutral market distribution at time T , and use the distribution to price European options. Details about the FFT application on option pricing can be found in Carr and Madan(1999).