

On the Relative Pricing of long Maturity Index Options and Collateralized Debt Obligations ¹

Pierre Collin-Dufresne²

Robert S. Goldstein³

Fan Yang⁴

First Version: October 2008
This Version: September 12, 2011

¹Previously entitled, "On the Relative Pricing of long Maturity S&P 500 Index Options and CDX Tranches". We thank Luca Benzoni, Josh Coval, Mike Johannes, Jakub Jurek, Erik Stafford, and seminar participants at Carnegie Mellon University, UT Dallas, HEC Lausanne, Federal Reserve Bank Washington DC, Madison-Wisconsin, Princeton, UCLA, EPFL, University of Vienna, Vanderbilt, ETH Zurich, Warwick, AQR, Bachelier World Symposium, BlackRock Quantitative Finance, Fields Institute, Kepos Capital, Moody's-NYU Credit conference, SwissQuote Conference, S&P Credit Risk Summit and especially Adrien Vesval for many helpful comments and Credit Suisse for generously providing the options data. All remaining errors are our own.

²Carson Family Professor of Finance, Columbia University and NBER, pc2415@columbia.edu

³C. Arthur Williams Professor of Insurance, University of Minnesota, and NBER, gold144@umn.edu

⁴Assistant Professor, School of Economics and Finance, The University of Hong Kong, fanyang@hku.hk

This is the pre-peer-reviewed version of the following article: Collin-Dufresne, Pierre, Robert Goldstein, and Fan Yang. "On the Relative Pricing of Long Maturity Index Options and Collateralized Debt Obligations." *Journal of Finance* 67, no. 6 (December 2012): 1983-2014, which has been published in final form at <http://dx.doi.org/10.1111/j.1540-6261.2012.01779.x>

On the Relative Pricing of long Maturity Index Options and Collateralized Debt Obligations

Abstract

We investigate a structural model of market and firm-level dynamics in order to jointly price long-dated S&P 500 index options and CDO tranches of corporate debt. We identify market dynamics from index option prices, and idiosyncratic dynamics from the term structure of credit spreads. We find that all tranches can be well priced out-of-sample before the crisis. During the crisis, however, our model can capture senior tranche prices only if we allow for the possibility of a catastrophic jump. Thus, senior tranches are non-redundant assets that provide a unique window into the pricing of catastrophic risk.

1 Introduction

Securitization – the process of pooling risky bonds and repackaging their cash flows to create new types of securities (‘tranches’) with varying levels of subordination – has been at the heart of the recent financial crisis. Indeed, the crisis began with the demise of two mortgage hedge funds run by Bear Stearns in 2007 due to their large bets in securitized portfolios of subprime mortgage bonds. It has been widely argued that the risks of subprime mortgages had been dramatically underestimated by market participants for a variety of reasons, not the least being that this was a relatively new market.¹ What is perhaps less well-known, however, is that the dramatic drop in prices of securitized portfolios was not limited to subprime collateral, but instead widespread across most major asset classes, such as commercial real estate, credit card receivables, senior unsecured loans, investment grade and high yield corporate debt.² Since availability of historical data on default experience is less of an issue for these other asset classes, many observers began to question whether there was a significant flaw in the pricing methodology used by Wall-Street to evaluate the prices of these securitized products. For example, a widely quoted article appearing in the popular financial press³ suggested that securitized products were mispriced because of the widespread use (in spite of its many known flaws) of the ‘Gaussian Copula Model’ (Vasicek (1987), Li (2000)) by Wall-Street firms.

The question of whether securitized assets were mispriced prior to the crisis has also been examined recently by the academic literature. For example, Coval, Jurek and Stafford (CJS, 2009) investigate the pricing of CDO tranches created from portfolios of investment-grade corporate bonds. In particular, they focus on the CDX.IG market, which is the most liquid of the CDO markets.

To price tranche spreads on the CDX.IG index, it is crucial to accurately identify both the systematic component and the idiosyncratic component of returns for the individual securities that compose the underlying portfolio. This is because the prices of these tranches are very sensitive to the level of correlation in returns across securities, and these correlations are increasing in the fraction of return volatility due to common movements and decreasing in

¹Data on subprime mortgage default history was scarce, see e.g., Hull and White (2010), Brunnermeier (2009).

²Investment-grade corporate credit spreads as measured by the CDX.IG index rose from 50bps in early 2007 to more than 250bps at the end of 2008. Even at the safest end of the spectrum the widening was dramatic. AAA-rated synthetic debt products, which were deemed virtually risk-free before the crisis, saw their spreads widen dramatically: CDX.IG super senior tranche widened from 5bps to 100bps, CMBX AAA “super duper” widened from 2bps to 700bps, ABS-HEL AAA tranche price rose from 0 to 20% upfront plus 500bps running. These numbers illustrate that it became much more expensive to insure AAA-rated debt across various markets (corporate, residential and commercial real estate).

³‘Recipe for Disaster: The Formula That Killed Wall Street,’ on Feb 23, 2009.

the fraction of return volatility driven by idiosyncratic movements. Following the insights of Breeden and Litzenberger (1978), CJS extract market dynamics from the prices of five-year index options across multiple strikes. They then assume that the idiosyncratic component is normally distributed, and calibrate the idiosyncratic volatility from observed equity returns. Finally, in the spirit of Merton’s (1974) structural model of default, CJS assume that each bond in the collateral portfolio defaults at its five year maturity if firm value falls below the default barrier. This default boundary is set so that the model matches perfectly the five-year credit spread on the collateral.

Their main findings are that market prices of senior tranches are too high (equivalently, senior tranche credit spreads are too low) and market prices of junior tranches are too low (junior spreads are too high) relative to those predicted by their model.⁴ Since their results appear to be 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 purchased senior claims solely with regard to their credit rating (a measure of expected default probability), but ignored the large systematic risk associated with holding a senior claim written on an underlying diversified portfolio.

This interpretation of their findings is rather puzzling, since participants in the market that CJS investigate tend to be rather sophisticated, and are active in both index options and CDS markets. Indeed, this market is the most liquid of the CDO markets, and is where hedge funds and major investment banks trade very actively. Further, CJS’s findings also imply that junior tranches – those with lowest rating and highest yield – which are typically purchased by hedge funds and other speculative investors (i.e., a different clientele than the ‘naive’ investors interested in the safe tranches), were severely underpriced, which raises further questions.

In this paper, we also investigate the relative prices of S&P 500 index options and spreads on the CDX index and synthetic CDO tranches based on the CDX index. An important difference between our framework and that of CJS is that we specify a *dynamic* structural model of default for each firm in the underlying portfolio. In particular, our model provides state-prices for all maturities, not just at the five year maturity. This allows us to jointly and consistently price stock options and CDO tranches 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

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.

take into account differences in the timing of defaults, which can significantly impact the cash flows of CDO 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 (risk-neutral) expected defaults and the specification of the idiosyncratic component of firm value dynamics. Using this calibration approach, we match the time series of tranche spreads well, in contrast to the results of CJS.

To gain some intuition for why it is essential to calibrate the model to match the term structure of CDX index spreads, we recall a well-documented failure of (diffusion-based) structural models of default: they dramatically underpredict default rates (Leland (2004)) and credit spreads (Jones, Mason, and Rosenfeld (1984), Collin-Dufresne and Goldstein (2001), Eom, Helwege, and Huang (2004)) for investment grade debt at short maturities. To match short-term spreads on investment grade debt, it is necessary to include jumps when specifying firm value 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, and in turn, the systematic risk (jump and diffusion components) of firm dynamics, the term structure of CDX index levels is useful for identifying total risk (jump and diffusion) of firm dynamics. Using these two pieces of information together allows us to identify the idiosyncratic component of firm 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 adjusted) expected losses (i.e., defaults) at shorter horizons. Without this calibration approach, defaults are *backloaded*, that is, occur later than actual market expectations. As it turns out, the timing of defaults is especially crucial for the equity tranche, which can be entirely wiped out by a few defaults.

The second important way our calibration approach impacts tranche spreads is by increasing the proportion of risk that is idiosyncratic versus systematic. Intuitively, as the proportion of systematic risk increases, the loss distribution becomes more fat-tailed. (Indeed, if there were no systematic risk and the portfolio was very large and homogenous, then, by the law of large numbers, the distribution would converge to a single point centered at the risk-neutral

expected loss – Vasicek (1987)).⁵ Below we show that if the model is not calibrated to match the term structure of CDX index spreads, then there is too little idiosyncratic jump risk, and thus too much loss probability mass is shifted into the tails. This implies that too much of the loss distribution is pushed into senior tranches, making their expected insurance payoff higher. This, in turn, tends to bias upward the break-even tranche spreads of senior tranches and downwards the equity tranche spread.

After calibrating the systematic risk component of the collateral portfolio to short and long maturity index options and the idiosyncratic risk component to the term structure of credit spreads, we find little evidence of relative mispricing between stock index options and CDO tranches before the crisis. In particular, senior tranche spreads (which CJS refer to as spreads on ‘economic catastrophe bonds’) seem in line with those predicted by our model.

During the crisis, however, we find that senior tranche spreads widen dramatically, in a way that can be reconciled within our structural model only by allowing for the possibility of a catastrophic jump. In that sense CDO tranches are not redundant securities. Instead, they provide a unique window into the market’s crash-risk expectation and risk-aversion. This is because the strike prices of traded options do not span far enough out of the money 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 (which collects a premium only when realized losses in the collateral portfolio exceed 30% - a type of event not witnessed even during the great depression). Therefore, during the crisis, we also investigate the performance of the model when we calibrate catastrophe-risk to match the super-senior tranche, and price all other tranches out-of-sample.

There is a large and increasing literature on correlated defaults.⁶ On CDO tranche pricing, Mortensen (2006), Longstaff and Rajan (2008), and Eckner (2009), have already demonstrated the ability of a few in-sample state variables to match the cross section of tranche spreads well. Our contribution with respect to the literature is to investigate the relative pricing across the stock option and CDO markets, specifically pricing CDO tranches out-of-sample.⁷

⁵We emphasize that this is a fairly ‘loose’ intuition, since it is not only the fraction of idiosyncratic to systematic risk that matters, but also the fraction of idiosyncratic risk that is diffusive versus jump risk. With more idiosyncratic risk coming from jumps as opposed to diffusion risk, we find a less fat tailed loss distribution.

⁶Examples 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, Freed, Geng, and Kapadia (2006), Das, Duffie, Kapadia, and Saita (2007), Duffie, Saita and Wang (2007), Jorion and Zhang (2009).

⁷Cremers, Driessen and Maenhout (2008) show option prices on individual firms are consistent with prices from the CDS market.

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 model predictions for time series of three and five year tranche prices relative to index options. In Section 5, we investigate the robustness of our results. We conclude in Section 6.

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

This paper investigates the relative pricing of CDO tranches and S&P 500 options within a structural framework. This requires a model of both market returns and individual firm level returns. In this section, we introduce a standard affine option pricing model for the market return, which we use for our benchmark. In particular, we extend the “preferred” stochastic volatility jump diffusion model of Broadie, Chernov and Johannes (BCJ, 2007) to incorporate two stochastic volatility processes. This additional state variable enables us to price both long and short term options consistently (BCJ use only one factor because they focus on short-dated options only).⁸

We then introduce a CAPM-like structural model for firm level returns that allows us to price the underlying portfolio of corporate bonds (or, for synthetic CDOs, the underlying portfolio of CDS), as well as the CDO tranches, which are essentially spread-options on the underlying CDO portfolio, as we make precise below.

2.1 Market dynamics for pricing SP500 options

There is a long tradition, since Breeden and Litzenberger (1978), to extract implied state price densities from quoted option prices.⁹ One common approach is to use a ‘local volatility’ model, which specifies a flexible implied volatility function (assumed to be deterministic in the underlying) which is then calibrated to match observed option prices across strikes and maturity.¹⁰ 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

⁸It is well-known that long-dated option implied volatility tend to be steeper than standard models, fitted to short term options would predict. (e.g., Das and Sundaram (1999), Backus, Foresi and Wu (2004), Carr and Wu (2003)). Christoffersen, Heston, Jacobs (2009) also find that two stochastic volatility factors is useful when using long dated options. However, they do not incorporate jumps in their analysis.

⁹See also Rubinstein (1994), Dupire (1994), Derman and Kani (1994) for early papers on extracting state prices from options.

¹⁰See, for example, Derman and Kani (1994), Dupire (1994) Rubinstein (1994), Dumas, Fleming and Whaley (1998).

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 extend far beyond the range of strikes of quoted option prices. Therefore, we prefer to use a parametric dynamic model. In particular, we specify risk-neutral market returns, and their associated variance state variables, to follow a joint-Markov affine jump-diffusion process. This approach guarantees that our ‘extrapolated’ volatility surface is both consistent (i.e., arbitrage-free) and generated from a model with economically motivated dynamics, with a relatively small number of parameters. One nice feature (which is necessary for our purpose) of having a dynamic option pricing model is that, once calibrated, we can obtain the state price density for all strikes and all maturities.¹¹

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 and Polson (2003), Eraker (2004), and Anderson and Andreasen (2000).¹² The recent paper by Broadie, Chernov and Johannes (BCJ, 2007) presents an encompassing test of various specifications proposed in the literature using an extensive data set on short-maturity S&P 500 futures options. Their preferred “SVCJ model” allows for stochastic volatility and correlated jumps in both stock returns and volatility processes. We therefore choose that model, but allow for two stochastic volatility factors. BCJ and indeed most of the literature focuses on short-maturity (i.e., less than 6 months) options. 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 it is possible to obtain longer dated options from the over-the-counter (OTC) market. We obtained a proprietary data-set of one- to five-year options from Credit Suisse to perform our calibration.¹³

Specifically, we define M_t as the value of the market portfolio, and (V_t, θ_t) as the two variance state variables. We specify their joint risk-neutral dynamics as:

$$\frac{dM_t}{M_t} = (r - \delta) dt + \sqrt{V_t} dw_1^Q + \sqrt{\theta_t} dw_2^Q + (e^y - 1) dq - \bar{\mu}_y \lambda^Q dt + (e^{y_C} - 1) (dq_C - \lambda_C^Q dt) \quad (1)$$

$$dV_t = \kappa_V (\bar{V} - V_t) dt + \sigma_V \sqrt{V_t} \left(\rho_1 dw_1^Q + \sqrt{1 - \rho_1^2} dw_3^Q \right) + y_V dq \quad (2)$$

¹¹In contrast, CJS use a local volatility model to fit only five-year options, since their model is essentially a static-one-period model.

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.

¹³In an earlier version of the paper, we used the five-year OTC dataset obtained from Citigroup by CJS (available on the AER website). Our new data is more comprehensive, and provides for an interesting cross-validation of the OTC data.

$$d\theta_t = \kappa_\theta(\bar{\theta} - \theta_t) dt + \sigma_\theta \sqrt{\theta_t} \left(\rho_2 dw_2^Q + \sqrt{1 - \rho_2^2} dw_4^Q \right) + y_\theta dq. \quad (3)$$

Here, dw_j^Q ($j = 1, 2, 3, 4$) are independent Brownian motions, and dq is a jump process with a constant jump intensity λ^Q . The market (M_t), and the variance state variables (V_t, θ_t) jump contemporaneously. The jump sizes of the variance state variables have exponential distributions, $y_V \sim \exp(1/\mu_V)$ and $y_\theta \sim \exp(1/\mu_\theta)$. The jump size in returns is assumed to be normal $y \sim N(\mu_y, \sigma_y)$.¹⁴ The compensator for the jump in the market price is

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

This is the standard SVCJ model studied in the literature (e.g., BCJ (2007)) extended for a second stochastic variance state variable with a correlated jump. We therefore refer to this model as the SV2CJ model. We note that 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 ‘catastrophic’ jumps for the pricing of senior tranches during the crisis. The impact of such jumps on equity prices and their expected returns has been investigated by Rietz (1988), and more recently revisited in Barro (2006), Backus, Chernov and Martin (2011) and others. Below we show that the spread (i.e., insurance premium) paid by buyers of the protection leg of the super-senior tranche is extremely sensitive to the (risk-neutral) likelihood of such crashes. In contrast, we show that the prices of options for those strikes that are actually traded are rather insensitive to changes in this likelihood. That is, given only observed option prices, we cannot uniquely identify state prices associated with catastrophic states of nature. Therefore, we investigate two models of market dynamics below: The first model has no catastrophic state, and is calibrated to match option prices only. The second model has a catastrophic state, and is calibrated to match both option prices and super-senior tranche spreads.

Given its affine dynamics, the (log) index return process has an exponential affine characteristic function. Therefore, European option prices can be solved by applying the Fast Fourier Transformation (FFT) (Heston (1993), Carr and Madan (1999), Duffie, Pan and Singleton (2000)). The solution of the log market characteristic function is given in Appendix 7.1.

¹⁴In standard SVCJ model (as in Duffie, Pan, and Singleton (2000)) this market jump size is normal conditional on the variance jump size. BCJ note that this correlation is extremely hard to estimate because jumps are rare events. Therefore, we follow BCJ in assuming that jump sizes are independent for parsimony.

2.2 Firm dynamics and structural default model

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

$$\begin{aligned} \frac{dA_t}{A_t} + \delta_A dt - r dt &= \beta \left(\sqrt{V_t} dw_1^Q + \sqrt{\theta_t} dw_2^Q + (e^y - 1) dq - \bar{\mu}_y \lambda^Q dt \right) \\ &+ (e^{y_C} - 1) (dq_C - \lambda_C^Q dt) + \sigma_i dw_i + (e^{y_i} - 1) (dq_i - \lambda_i^Q dt). \end{aligned} \quad (5)$$

This is a CAPM-like equation for individual firm's asset return, where β denotes the loading of each firm's asset return dynamics on the market (excess) return. There are a few differences worth noting: First, unlike in the standard continuous time CAPM model, we allow the market to exhibit systematic jumps. Second, we allow for idiosyncratic jumps $dq_i(t)$ with a risk-neutral intensity whose deterministic dynamics are calibrated to match the term structure of CDX index spreads. Third, we specify β , which is the coefficient of the regression of the unlevered asset value on the levered market return (ignoring the common catastrophe risk) to be constant.¹⁵ Lastly, we note that 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.

Under these specifications, the log market and log asset values have dynamics

$$\begin{aligned} d \log M_t &= \left(r - \delta - \bar{\mu}_y \lambda^Q - (e^{y_C} - 1) \lambda_C^Q - \frac{1}{2} V_t - \frac{1}{2} \theta_t \right) dt + \sqrt{V_t} dw_1^Q + \sqrt{\theta_t} dw_2^Q \\ &+ y dq + y_C dq_C. \end{aligned} \quad (6)$$

$$\begin{aligned} d \log A_t &= \left(r - \delta_A - \frac{1}{2} \beta^2 V_t - \frac{1}{2} \beta^2 \theta_t - \frac{1}{2} \sigma_i^2 - \beta \bar{\mu}_y \lambda^Q - (e^{y_C} - 1) \lambda_C^Q - (e^{y_i} - 1) \lambda_i^Q \right) dt \\ &+ \beta \sqrt{V_t} dw_1^Q + \beta \sqrt{\theta_t} dw_2^Q + \sigma_i dw_i + \log [\beta (e^y - 1) + 1] dq + y_C dq_C + y_i dq_i. \end{aligned} \quad (7)$$

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:

$$\tau_i = \inf\{t : A_i(t) \leq A_B\}. \quad (8)$$

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

¹⁵We note that a perhaps preferable assumption might have been to specify the coefficient of the regression of unlevered firm value on unlevered market value to be constant. This would have complicated our computations and calibration, as the option on the levered market portfolio would not be as straightforward to calculate.

2.3 Basket CDS index

In our application, we use data on synthetic CDO tranches based on the Dow Jones CDX North American Investment Grade Index.¹⁶ The underlying portfolio is an equally weighted basket of 125 investment grade single name credit default swaps (CDS). It can essentially be thought of as a portfolio of 125 liquid five-year CDS, each with investment grade status. Recall that a single name CDS is an insurance contract by which a protection buyer makes quarterly payments of the CDS premium (equal to the CDS spread times the notional) to the protection seller for the entire maturity of the contract, unless the underlying firm experiences a default event. In the latter case, payments stop and the protection buyer delivers the defaulted bond to the seller in return for the notional (effectively receiving the loss-given default on the bond). The spread of a single-name CDS on a firm is closely related to the credit spread of a bond on the same firm (since buying protection and buying the bond essentially results in a default-free position). In fact, Duffie (1999) discusses conditions under which a pure arbitrage argument makes the two equal.¹⁷ The running spread on the CDX index is closely related to a weighted average of single name CDS spreads. The protection buyer on the CDX contract makes quarterly payments (and sometimes an upfront payment) on the outstanding notional of the contract until the maturity of the contract, or until one of the constituents experiences a default event. In the latter case, the protection buyer delivers a defaulted bond for the notional amount in the index to the protection seller in return for the notional (1/125 of the index notional) of the bond in the index. Essentially, then, the buyer receives the loss given default on the defaulted entity in the index. The CDX index notional is then reduced by 1/125, and the contract continues.

To determine the index spread,¹⁸ the present value of cash flows that go to the protection

¹⁶The CDX.IG index is a *synthetic* CDO market in that the underlying collateral is a portfolio of credit default swaps (CDS), rather than a portfolio of bonds, as would be the case for a more traditional cash CDO. Thus, this market is *unfunded*, which implies that trades require smaller margins, and therefore offer more leverage, and in turn, more liquidity. The prices on the tranches and on the collateral are rather transparent, as quotes for all tranches are sent from all dealers to market participants throughout the day, and large notionals can be executed at fairly small bid/ask spreads (at least before the crisis). Adding to the transparency of this market, prices on all the CDS that constitute the collateral are also readily available to all market participants and, for the vast majority, quite liquid. In this respect, the CDX.IG market is very different from the subprime CDO market, which has been described by many as opaque (e.g., Gorton (2010)).

¹⁷Empirically, the difference between CDS and corporate bond credit spreads (i.e., the ‘CDS basis’) is close to zero or slightly positive when the risk-free benchmark rate used to compute the bond spread is based on the LIBOR-swap curve (Blanco, Brennan and Marsh (2005)). During the 2007-2009 crisis, however, the basis became very negative, indicating this arbitrage relation broke down (Bai and Collin-Dufresne (2011)).

¹⁸There are some technical differences between the CDX contract specification and a portfolio of CDS. Most notably, the CDX is partially settled in upfront based on a fixed running coupon spread. For our exposition here we focus on a full running (i.e., zero upfront) CDX index specification (see Collin-Dufresne (2009) for further discussion of the institutional details).

buyer (the ‘protection leg’) and protection seller (the ‘premium leg’) are set equal to each other. The values of these two cash-flow ‘legs’ are obtained by computing the following expectations (assuming a one dollar total notional):

$$V_{idx,prem}(S) = S E^Q \left[\sum_{m=1}^M e^{-rt_m} (1 - n_{t_m}) \Delta + \int_{t_{m-1}}^{t_m} e^{-ru} (u - t_{m-1}) dn_u \right] \quad (9)$$

$$V_{idx,prot} = E^Q \left[\int_0^T e^{-rt} dL_t \right]. \quad (10)$$

Here, we have defined the cumulative notional of defaulted firms in the portfolio by $n_t = \frac{1}{N} \sum_i \mathbf{1}_{\{\tau_i \leq t\}}$, and the cumulative loss in the portfolio as:

$$L_t = \frac{1}{N} \sum_i \mathbf{1}_{\{\tau_i \leq t\}} [1 - R_i(\tau_i)], \quad (11)$$

where $R_i(t)$ is the recovery rate on firm i when it defaults at time t . Expectations are all taken with respect to the risk-neutral Q-measure, which we assume exists, to rule out arbitrage (Harrison and Kreps (1979)).

2.4 CDO tranche spreads

For each maturity that the CDX index is traded, there are also six CDO tranches based on that index that are traded. Each tranche is characterized by its ‘attachment points,’ which are related to its subordination level. For the CDX.IG underlying portfolio, 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 of a particular L-U% tranche makes periodic premium payments (corresponding to the remaining tranche notional times the tranche spread) until the contract expires. In return she receives protection payments if cumulative losses in the underlying CDX index exceed L%. These payments are made as default events occur and are equal to the realized losses in the remaining underlying portfolio. Payments stop when cumulative losses in the underlying portfolio exceed U%, after which the tranche notional is exhausted, and the contract ends. 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 the cumulative losses (L_t) in the portfolio underlying the tranche is:¹⁹

¹⁹To simplify notation, we assume that the underlying portfolio has a 1 dollar total notional, and we price each tranche as a claim to a notional corresponding to its ‘natural’ width. Specifically, we assume the equity tranche has a notional of 3%, the mezz a notional of 4%, . . . , and the super-senior a notional of 70%. In practice, tranche upfront and spread quotes are given in percentage of the chosen notional protection purchased on the specific, say L-U%, tranche. The appropriate underlying portfolio notional is obtained by dividing the tranche notional by its width (U-L).

$$\begin{aligned}
T_j(L_t) \equiv T_{K_{j-1}, K_j}(L_t) &= \max[\min(L_t, K_j) - K_{j-1}, 0] \\
&= \max[L_t - K_{j-1}, 0] - \max[L_t - K_j, 0].
\end{aligned} \tag{12}$$

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

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

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

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

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

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

The reason, we suspect, for these two different practices is that the pricing of the equity tranche using equation (14) is, as we show below, 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. We interpret the fact that the market has different cash-flow mechanisms for equity tranche versus the other tranches as evidence that market participants recognize the sensitivity of equity tranches to the timing of default and therefore, no doubt, also would adjust their models to account for this. This is, in part, why we think a dynamic model is useful.

Finally, the super-senior tranche premium is specified by

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

The term $n_u R$ in the integrand captures the fact that premium payments are reduced by the amount recovered in default.²⁰ We first focus on the pre-crises period September 2004 to September 2007. We then investigate how the model performs during the crisis.

²⁰Theoretically, if recovery levels exceed 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 this possibility, however.

3 Data and Calibration

We calibrate the model parameters and identify the time series of state variables in two steps. First, we calibrate market return parameters and state variables using data from over-the-counter long- and short-dated European option prices. Then, we calibrate firm-specific parameters (distance to default, idiosyncratic diffusion and jump risk parameters, and market-beta) using firm-specific information from CRSP Compustat as well as credit spread information. We first present the data used, and then explain our methodology.

3.1 Data

Our primary data includes one- and five-year S&P 500 European option implied volatilities, the CDX North American Investment-Grade Index spreads from one to five years, and tranche spreads written on this index for three- and five-year maturities. The option data are from Credit Suisse. For each date and maturity, there are eleven option implied volatilities with moneyness from 0.5 to 1.5 with 0.1 increments.²¹ The CDX index and tranche data are from JP Morgan. Every six months (on March 21 and September 21) a new on-the-run CDX series is introduced. The composition of the CDX index may also be changed at that time (constituents can be removed or added if there has been a default or a downgrade, for example). The contracts based on old (off-the-run) series continue to trade, but become less liquid as the maturity of the contract shortens. Many participants choose to roll into the newest series (the on-the-run) contract. For our analysis we distinguish two sub-periods: the ‘pre-crisis’ period (September 21, 2004 to September 20, 2007) and the ‘crisis’ period (September 21, 2007 to September 20, 2008). The pre-crisis period includes data from on-the-run Series 3 through Series 8, whereas our crisis period includes data from on-the-run Series 9 and Series 10. To simplify the calibration and reduce microstructure-related issues (such as bid-ask bounce), we use only the closing quotes of every fourth Wednesday for our calibration procedure. Thus, for every series, our calibration data consists of seven observation dates of i) option implied volatilities, ii) the CDX index, and iii) tranche spreads.

To identify parameters such as asset betas, idiosyncratic volatilities, payout ratios and leverage ratios associated with firm value dynamics, we also use daily CRSP and quarterly/annual

²¹Compared to CJS, who use five-year OTC option data from Citigroup, our deep out-of-the-money options have slightly higher implied volatilities on average. The price difference of the same options from different sources indicate large bid-ask spreads for these long-term options. More importantly, our option data imply higher risk-neutral market downside risk than CJS. Therefore, calibrating to Credit Suisse data should increase the mis-pricing documented by CJS in that it would push senior tranche spreads higher relative to equity tranche spreads (relative to a calibration using option data from Citigroup).

Compustat data. The composition of every CDX collateral pool are from Markit. Swap rates are from Datastream. Dividend yields are from OptionMetrics.

3.2 Calibration of market dynamics

Since market dynamics are specified to be affine, standard results (e.g., Heston (1993), Duffie, Pan, and Singleton (2000)) allow us to compute option prices in closed-form (up to a Fourier Transform, see Appendix 7.1). We use this closed-form expression to minimize the relative Root Mean Square Error (RMSE) between model prices and observed prices by searching over both parameters and latent state variables (θ_t, V_t) . From the point of view of our procedure, the only difference between state variables and parameters is that state variables change each observation date, whereas parameters remain constant over each six month period associated with a given series. This recalibration every six months allows us to identify regime shifts.

We also investigate two different specifications for the market model (and a few more in the robustness section below). For the first specification, we ignore catastrophic jump risk (i.e., we set the catastrophic jump intensity $\lambda_C^Q = 0$) and calibrate the market dynamics using only option data. For the second specification, we allow for a catastrophic jump of size ($y_C = -2$), which is associated with a -87% drop in price, and calibrate a piece-wise linear risk-neutral intensity to match both the three- and five-year super-senior tranche spreads, as well as option prices. We do this because our results below suggest that, especially during the crisis, quoted option prices do not contain enough information about the very deep out-of-the-money state prices, and these are crucial for the pricing of senior tranches. The parameter values are reported in Table 1.

In the pre-crisis period, we report only one set of parameters for each series (obtained for $\lambda_C^Q = 0$), because the implied catastrophic jump intensities calibrated from super-senior tranche spreads are very close to zero, so option parameters do not change much. In contrast, during the crisis-period λ_C^Q becomes much more sizable, and therefore the two model specifications require significantly different model parameters.

In Appendix 7.2 we discuss our calibration results in more detail. We show that the fitted state variables seem sensible, in that the implied volatility tracks the VIX fairly closely (Figure 10). Further, we show that the crash intensities are close to zero during the pre-crisis period (except for a short spike during May 2005, which corresponds to the downgrade of Ford and General Motors), but spike up to 100 bps for the one- to three-year intensity during the crisis (see Figure 11).

A few parameter estimates in Table 1 are worth mentioning. In particular, we see that the

two volatility state variables have very different persistence parameters, with θ_t much more persistent than V_t . Therefore, θ_t has greater impact on long-term options, whereas V_t has greater impact on short-term options. This explains why two state variables are useful in explaining the term structure of options. While the parameters are not exactly the same as in BCJ (as we have a slightly different model specification), they tend to be similar. We find that options imply a ‘standard’ market jump (μ_y) of approximately -40%, with risk-neutral frequency of once every ten years in normal times.

Figures 1 and 2 show the fit of the model for our calibration as well as the corresponding implied risk-neutral distribution during the pre-crisis and crisis period respectively. It is apparent that both models do a good job at fitting the sample of 11 long-dated implied volatilities during both periods. We emphasize that the models also fit short dated (one year) volatilities well.

Interestingly, the figures also show that both calibrations (with and without catastrophe jump) imply almost identical option prices for those strikes that are actually traded, that is, for option strikes with moneyness greater than 0.5. Indeed, calibrating market dynamics to match the super-senior tranche impacts state prices significantly only for moneyness levels around 0.2 or lower in general (as seen in panel B). During the crisis there are more marked differences between the two models’ risk-neutral densities. This is our first indication that one cannot ‘extrapolate’ the information in option prices to deduce information regarding the super-senior tranche.

3.3 Calibration of Firm value dynamics

After estimating the parameters and state variables for market dynamics using the method described above, we then estimate the firm-specific parameters of the asset dynamics in equation (5). These parameters include: payout ratio (δ_A), asset-market beta (β), idiosyncratic diffusion (σ_i), jump risk intensity (λ_i) and size y_i , default boundary (A_B), and the loss given default (ℓ).

For our benchmark calibration we make the assumption that the 125 firms in the CDO collateral portfolio are homogeneous and therefore calibrate these parameters to average firm characteristics. (In our robustness section below, we investigate the impact of firm heterogeneity.) We choose the parameters to match various statistics constructed from individual firm level data from CRSP Compustat for the specific constituents of each CDX index series. In particular, we estimate, for each series, the average asset beta, idiosyncratic asset volatility, payout ratio and leverage ratio. These numbers are reported in Table 2 below for each series

Parameter	Pre-crisis										Crisis					
											$\lambda_c = 0$			$\lambda_c > 0$		
	Series 3	Series 4	Series 5	Series 6	Series 7	Series 8	Series 9	Series 10	Series 9	Series 10	Series 9	Series 10	Series 9	Series 10		
κ_V	4.316	4.800	4.836	3.980	2.178	0.877	4.886	5.001	4.323	4.815						
\bar{V}	0.0018	0.0042	0.0046	0.0054	0.0057	0.0036	0.0015	0.0015	0.0012	0.0013						
σ_V	0.2961	0.274	0.2732	0.2666	0.2422	0.3296	0.2578	0.2613	0.2715	0.2531						
ρ_1	-0.48	-0.48	-0.48	-0.48	-0.48	-0.48	-0.48	-0.48	-0.48	-0.48						
μ_V	0.0503	0.0504	0.0491	0.0425	0.0736	0.0284	0.046	0.0458	0.0844	0.0618						
κ_θ	0.00130	0.0012	0.0012	0.0015	0.0015	0.00050	0.0012	0.0012	0.0012	0.0012						
$\bar{\theta}$	0.0068	0.0056	0.0057	0.0044	0.0055	0.0057	0.0044	0.004	0.0041	0.0049						
σ_θ	0.00068	0.00074	0.00075	0.00075	0.00075	0.00069	0.00080	0.00081	0.00073	0.00072						
ρ_2	0.00032	0.00034	0.00034	0.00033	0.00027	0.00039	0.00035	0.00035	0.00036	0.00036						
μ_θ	0.0668	0.0668	0.0667	0.0484	0.0281	0.0208	0.0647	0.0652	0.0221	0.0355						
μ_y	-0.3816	-0.3834	-0.3796	-0.5038	-0.2883	-0.4723	-0.4439	-0.4415	-0.4584	-0.4369						
σ_y	0.0167	0.0173	0.0171	0.0177	0.0205	0.0231	0.0177	0.0178	0.0175	0.0171						
λ	0.0886	0.1089	0.1179	0.0847	0.1598	0.0991	0.1726	0.1828	0.192	0.1496						
RMSE ($\lambda_c = 0$)	2.27%	1.01%	0.92%	1.28%	0.78%	2.77%	1.94%	0.93%	1.86%	1.74%						
RMSE	2.27%	1.68%	0.92%	1.38%	0.82%	2.30%										

Table 1: Calibrated parameters of market dynamics as described in Section 3.2.

(as well as the risk-free rate estimate we use). We explain in greater detail in Appendix 7.3 how we construct these key statistics from the raw CRSP-Compustat data (our method essentially follows Schaefer and Strebulaev (2008) closely).

Consistent with recent literature, we set the default boundary (A_B) to equal 60% of the leverage ratio reported in Table 2. 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 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% of face value, 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), Davydenko, Strebulaev and Zhao (2011)). 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 (2010), who estimates this ratio to be in the range $\frac{A_B}{F} \in (56\%, 70\%)$.

We set the (risk-neutral) recovery rate in normal times to be 40%, 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, Brady, Resti, and Sironi (2005)). As a recent example, the recovery rate on General Motors debt, whose default occurred during the crisis, was approximately 10%. Therefore, in our model we set recovery rate to 20% in the event that a catastrophic market jump occurs.

The most crucial distinction between our calibration procedure and that of CJS is the estimation of idiosyncratic dynamics. CJS assume that idiosyncratic risk is driven only by diffusion risks, 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 (we also match the diffusion component of idiosyncratic risk to the time series estimate of idiosyncratic risk). Specifically, we set the idiosyncratic jump size to $y_i = -2.0$ (a value

Series	Period	Asset beta	Idiosyncratic asset volatility	Leverage ratio	Payout ratio	Risk-free rate	S&P 500 div. yield
3	9/2004-3/2005	0.56	19.2	37.3	1.87	1.81	1.64
4	3/2005-9/2005	0.57	18.7	36.3	2.37	2.88	1.78
5	9/2005-3/2006	0.60	19.0	33.4	2.73	3.90	1.92
6	3/2006-9/2006	0.61	18.9	32.9	3.04	4.75	2.00
7	9/2006-3/2007	0.62	19.1	32.2	3.08	5.08	1.95
8	3/2007-9/2007	0.61	18.8	31.7	3.06	4.83	2.00
9	9/2007-3/2008	0.64	18.4	30.6	2.64	2.95	2.08
10	3/2008-9/2008	0.66	17.9	28.8	2.43	1.87	2.14

Table 2: Key aggregate statistics of collateral firms, risk-free rate, and S&P 500 index dividend yield for six-month period that a given series is on-the-run. Estimates other than asset beta are annualized and reported in percentage terms.

which basically guarantees default), and then calibrate the idiosyncratic jump intensity (which we assume to be a deterministic step function, as for the catastrophic jump intensity above) to match the term structure of credit spreads on the underlying CDX index (with maturities one to five years). We find that including idiosyncratic jumps is essential for matching short-maturity spreads. Indeed, it is well-documented that diffusion-based structural models of default for investment grade firms fail badly at capturing credit spreads (i.e, risk-neutral expected losses) at short maturities. Moreover, the systematic jump risk component (which was calibrated to match option prices) is grossly insufficient for capturing short dated credit spreads on individual credits. That is, our calibration suggests that short term credit spreads on the CDX index require a substantial amount of idiosyncratic jump risk. As we show below, calibrating the model to match the term structure of CDX index spreads has significant implications for the pricing of CDO tranches.

In Appendix 7.4 we give more details about our estimation procedure for the intensity parameters. We also plot the resulting intensity parameters in Figures 12 and 13 (for the two calibrations with and without catastrophic risk respectively²²). As one would expect from our discussion above, the one-year idiosyncratic jump intensity follows closely the one year credit spread. The relation is less tight at the longer end, where diffusion risk starts to play a bigger role in contributing to expected losses.

²²As one would expect calibrated idiosyncratic jump intensities are higher in the case without catastrophe risk, to compensate for the absence of market jump risk.

3.4 Discussion

As is clear from Table 2, the relative contribution of systematic and idiosyncratic risk to total risk shifted progressively during this period, with the fraction of total risk due to systematic risk increasing steadily as the crisis unfolded. This has immediate implications for the resulting credit loss distribution. We expect that, as the proportion of systematic risk increases, the loss distribution becomes more fat-tailed. (Indeed, if there were no systematic risk and the portfolio was very large and homogenous, then, by the law of large numbers, the distribution would converge to one point: the expected loss - Vasicek (1987)). This can be seen in Figure 3, which plots a representative loss distribution pre-crisis and during crisis (panel A), as well as the difference between the two (panel B). Clearly, the loss distribution has fatter tails during the crisis. This will have implications for the relative pricing of senior and junior tranches, as we show next.

4 Results

Given parameter estimates described in the previous section, we use our time series of long- and short-term index option prices to estimate, each week, the volatility state variables (which minimize the relative RMSE given our parameter choices explained in the previous section). Using our time series estimate of the initial firm leverage as given in Table 2 we can calculate tranche prices (and spreads) as explained in section 2 for all three and five year tranches. We first discuss the time series average results of these tranche spreads compared to the data for the pre-crisis and crisis period, and then look at the time series results in more detail.

4.1 Average tranche spreads results

In this section we report average tranche spreads for the three- and five-year tranches and for various sub-periods (pre-crisis period September 2004 - September 2007, crisis period September 2007-September 2010, and full sample). We compare the following specifications in Table 3:

- The actual values (historical data averages);
- Benchmark model: catastrophe 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 catastrophe jumps; Idiosyncratic jumps calibrated to match the 1-year, 2-year, 3-year, 4-year and 5-year CDX index spreads;

	1-year	2-year	3-year	4-year	5-year
Data	14	20	27	35	44
Benchmark	13	20	27	35	44
$\lambda_C = 0$	14	21	27	36	44
$\lambda_C = \lambda_i = 0$	1	7	18	31	44

Table 4: Historical and model-estimated average CDX index spreads September 2004 - September 2007 for four different models: i) benchmark, ii) benchmark without catastrophe jump ($\lambda_C = 0$), iii) benchmark without either catastrophe jump or idiosyncratic jumps ($\lambda_C = 0, \lambda_i = 0$).

- $\lambda_C^Q = 0, \lambda_i^Q = 0$: No catastrophe jumps; No idiosyncratic jumps; Default boundary calibrated to match only the 5Y CDX index;
- The results reported by CJS (which is available only for a sub-period).

As shown in the last row, our framework generates errors that are an order of magnitude smaller than those 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)), it is likely the fit could be improved even further if one were to calibrate market and firm dynamics by including information from other, say mezzanine, tranche spreads in our calibration, rather than pricing them out-of-sample, as we do here. But, the more interesting take away from our out-of-sample exercise is that there does not appear to be any blatant evidence for mispricing when we use our model specification to price tranches relative to index options, in stark contrast to the evidence of overpricing of ‘economic catastrophe bonds’ documented in CJS. To understand why our results are so different, we look at alternative specifications of the model, where we turn off various features that distinguish our model.

We begin with the case closest to CJS; the model with no catastrophe 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. Indeed, in Table 4 we show the CDX index spreads for maturities of 1-4 years predicted by each of the models. We see that the ($\lambda_C^Q = 0, \lambda_i^Q = 0$) model predicts credit spreads of 1bp and 7bps at maturities of 1 and 2 years, respectively, well below the market quotes of 14bp and 20bp. As such, expected losses are “backloaded” in this model, 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.

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 2.8%. Indeed, the standard deviation of the loss distribution without idiosyncratic jumps ($\lambda_C^Q = 0, \lambda_i^Q = 0$) has a standard deviation of 6.8%, whereas the standard deviation of the loss distribution with idiosyncratic jumps ($\lambda_C^Q = 0$) has a standard deviation of 2.7%. 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. And indeed, that model underpredicts the spreads on senior tranches. This problem can be solved by adding a catastrophe jump, since senior tranches tend to be mostly sensitive to these catastrophe events. As shown previously, including a catastrophe jump has virtually no impact on the model’s ability to match option prices. Our benchmark model then sets the catastrophe jump intensity so as to match the super senior tranche spread perfectly, and seems to do well pricing all other tranches ‘out-of-sample.’

In summary then, we find that in order to estimate tranche spreads, it is necessary that the model is 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/risk-neutral expected 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 (especially during the crisis). However, spreads on other tranches can be *interpolated* reasonably well given option prices and super-senior tranche spreads. We conclude that S&P 500 options and CDX tranche prices market can be fairly well reconciled within our arbitrage-free model, in contrast to the findings of CJS. In that sense these two markets appear to be well integrated.

Table 3 shows that these results hold for both the three- and five-year tranches and in both subperiods (i.e., during the pre-crisis as well as during the crisis period). However, the table also shows that our model has difficulty matching closely the tranche prices during the crisis period. Indeed, the benchmark model predicts too high equity tranches despite matching the super-senior perfectly. Our interpretation of these results is that the widening in credit spreads on the CDX.IG index, especially at the short end, was so large during the crisis that, despite the increase in systematic jump risk implied from the increase in implied volatilities on index options, it required a very large amount of idiosyncratic jump risk for our model to match short CDX term spreads. Indeed, in the Appendix we show that the implied idiosyncratic risk due to jumps increased dramatically during the crisis (see figures 12 and 14). Of course, during the crisis many market prices displayed rather surprising behavior so it is perhaps not surprising that the relative tranche prices appear to behave so differently than before the crisis.

4.2 Time-series of tranche spreads

In the previous section we analyzed the time series average of the predicted tranche prices. We showed that our benchmark model is able to fit average historical spreads across all tranches

well (especially pre-crisis), once it is calibrated to match the term structure of CDX index spreads and the super-senior claim. In this section we investigate the time-series performance.

Recall that 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 (their time series are shown in Figure 4). With this calibration, we then estimate tranche spreads. The results for the three-year tranches are given in Figures 5 and 6, and for the five year tranches in Figures 7 and 8, for the models with and without catastrophe risk, respectively. Consistent with the time-averaged results, the pictures reveal a dramatic improvement in fit relative to the model proposed by CJS. Note that the performance is somewhat better during the pre-crisis period than during 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. While the fit of the model could be further improved on these tranche spreads by including them during the calibration procedure, this is not our purpose here. Instead, we conclude that the overall fit of the model is quite good (especially pre-crisis) and certainly does not lend much credence to the notion that there is clear evidence for relative mis-pricing between the S&P 500 option and CDX tranche markets, let alone that there would be possible arbitrage opportunities.

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 (i.e., we offer no judgement about whether the absolute levels are ‘correct.’) One caveat is that we recalibrate the intensity parameters every week to fit the term structure of credit spreads. It would be interesting to study an extension where the term structures of idiosyncratic and catastrophe jump risks would be governed by stochastic processes whose dynamics could be inferred from tranche prices. Indeed, the time series results clearly suggest that these intensities are time-varying and stochastic. For example, it seems that pre-crisis the catastrophe intensity was close to zero, but that during the crisis the catastrophe intensity increased dramatically. This can be seen by comparing the benchmark model predicted super-senior tranches with and without catastrophe jumps in Figures 7 and 8 respectively.

Also, comparing the results for the three-year tranches given in Figures 5 and 6 we see that the model performs equally well for the three year tranches as for the the five-year tranches. We now turn to an analysis of the robustness of our findings, by relaxing several assumptions we made for our benchmark model specification.

5 Robustness

For parsimony, many simplifying assumptions were made in the previous sections, such as i) static capital structure, ii) two state variables driving market return variance, iii) constant interest rates, iv) uncorrelated idiosyncratic shocks (i.e. “no industry effects”), v) firm homogeneity, and vi) zero cash holding of the firms. Here we show that our results are robust to all these assumptions.

We investigate deviations from our benchmark model along several dimensions. For each deviation, we continue to calibrate the model to one- and five-year option implied volatilities, one- to five-year CDX indices, and the super-senior tranche spreads. Then we report model implied tranche spreads for three different representative days in our data set. We choose the median, 25th and 75th percentile of the five-year CDX index during the pre-crisis period (to provide some perspective on the state and sub-period dependence of our results). The results are given in Table 5 for the following extensions of our benchmark model:

- **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, Ju, and Leland (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’.
- **SVCJ:** To show that our results are not driven by our specific choice of option model, we also consider market dynamics as in the SVCJ option model of BCJ, and calibrate it to best match historical index option prices. Since there is one less state variable and a few less parameters, the SVCJ model fits the index options less well. However, the CDX tranche spreads predicted by the model are still very close to the benchmark.
- **Stochastic short-term rate:** We specify the spot rate to follow Vasicek (1977) interest rates and calibrate the model to the average term structure of treasury rates for maturities of three-month and one- to five-year. 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 Vasicek-type stochastic interest rate is also affine, so that we can apply 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”.

- Financials and industrials: We introduce heterogeneity across firms by distinguishing financials and industrials. In the pre-crisis period, there are about 23 financials and 102 industrials in the collateral pool. Market dynamics remain the same as in the benchmark case. We estimate aggregate firm parameters such as asset beta, idiosyncratic asset volatility, payout ratio, and leverage ratio separately for the two sectors. In particular, for financials, we add deposit, which earns the risk-free return, to their capital structure. We set deposit to be 80% of their asset values, and leverage up their firm parameters respectively. Then, we calibrate the idiosyncratic jump intensities of financials to their average CDS spreads with maturities from one to five years. CDS data are from Datasream. Their default boundaries are set to be 90% of the leverage ratio, since most of their debt is safe deposit. We calibrate a CDX model of two sectors, where 23 homogeneous financial firms are characterized as above and the idiosyncratic jump intensities of 102 homogeneous industrials are calibrated to the CDX indices.
- Accounting for cash holdings: Using Compustat data, we estimate the average cash holdings of firms in the CDX index to be approximately 5% across all series. Here we define the leverage ratio as (debt minus cash holdings) divided by total assets. We then assume that the default boundary is 60% of this net leverage. In order to match the term structure of CDS spreads, this reduction in the default boundary implies a larger estimate of idiosyncratic jump risk. Therefore, it further reduces the amount of weight in the tails, and in particular, in the senior tranches. But as we can see from the table, this adjustment does not change tranche spreads significantly from the benchmark.

Overall, these robustness checks imply that, while these different assumptions do impact tranche spreads to some degree, 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.

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 quite 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.

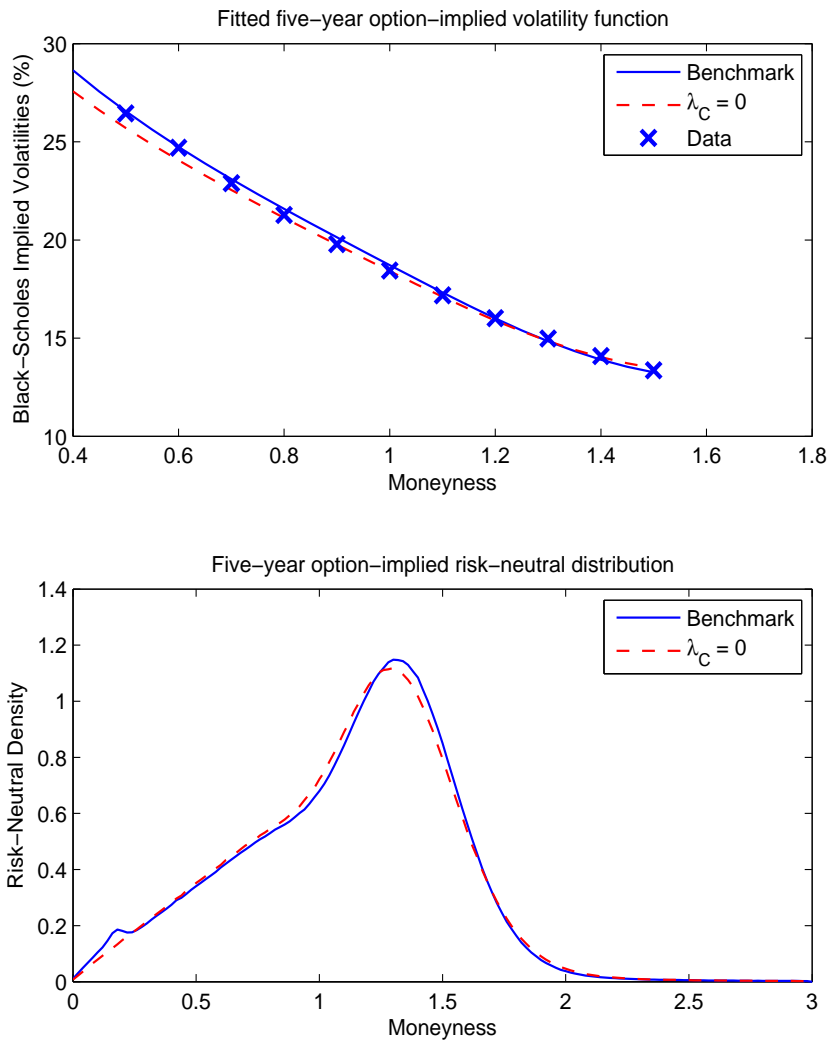


Figure 1: Market dynamics given in equations (1)-(3) are calibrated to match one and five year option prices on June 15, 2005 (the pre-crisis period). The parameters are given in Table 1. The upper panel shows the model-implied five year volatility surface and the actual option prices as a function of moneyness when there is a catastrophe jump (benchmark), and when there are none ($\lambda_C = 0$). The lower panel shows the corresponding risk-neutral distributions for the five-year market index level.

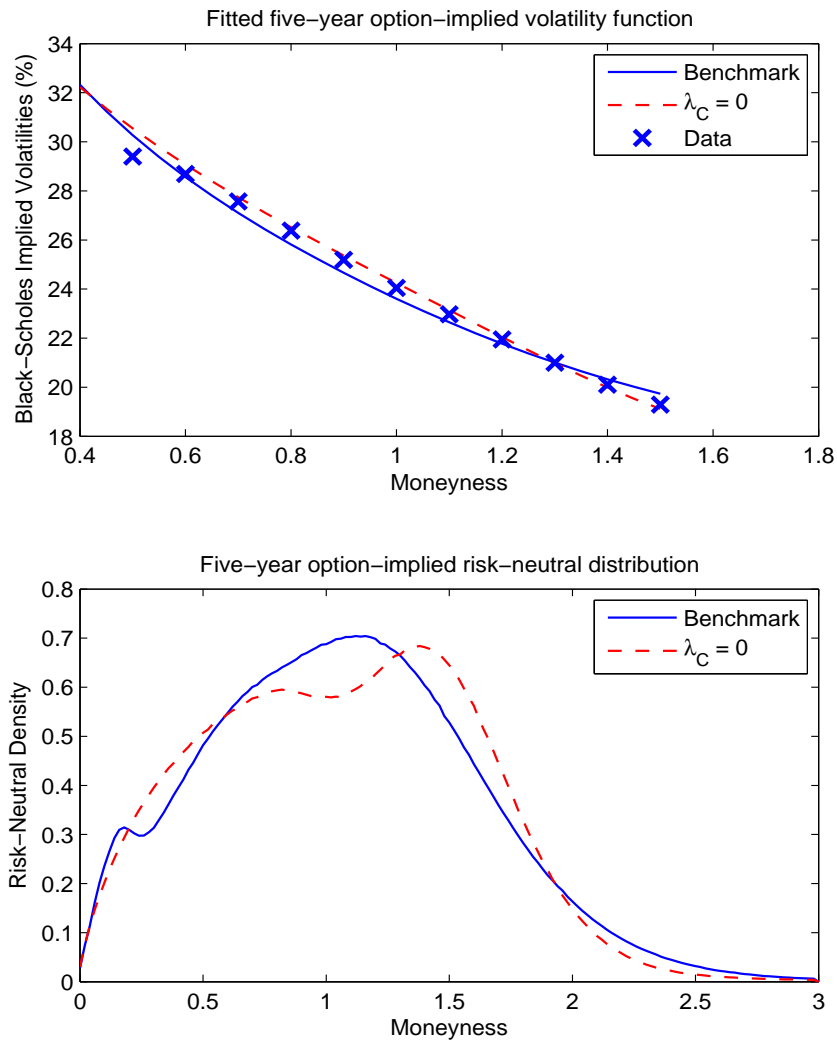


Figure 2: Market dynamics given in equations (1)-(3) are calibrated to match one and five year option prices on May 21, 2008 (the crisis period). The parameters are given in Table 1. The upper panel shows the model-implied five year volatility surface and the actual option prices as a function of moneyness when there is a catastrophe jump (benchmark), and when there are none ($\lambda_C = 0$). The lower panel shows the corresponding risk-neutral distributions for the five-year market index level.

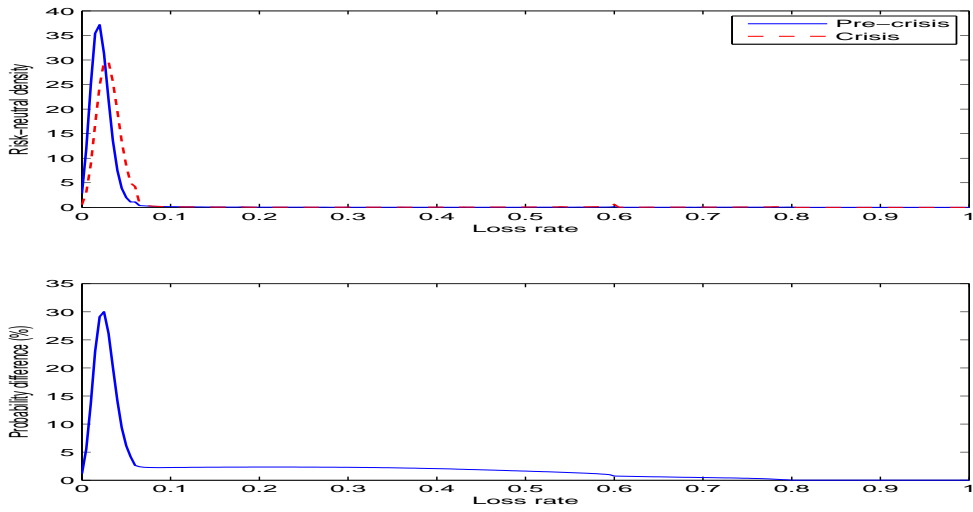


Figure 3: The upper panel shows the risk-neutral loss density for the pre-crisis and crisis periods. The crisis period has higher expected losses and a less-peaked distribution due to a larger proportion of risk being systematic. The lower panel shows the difference in the cumulative loss distributions for the crisis and pre-crisis periods.

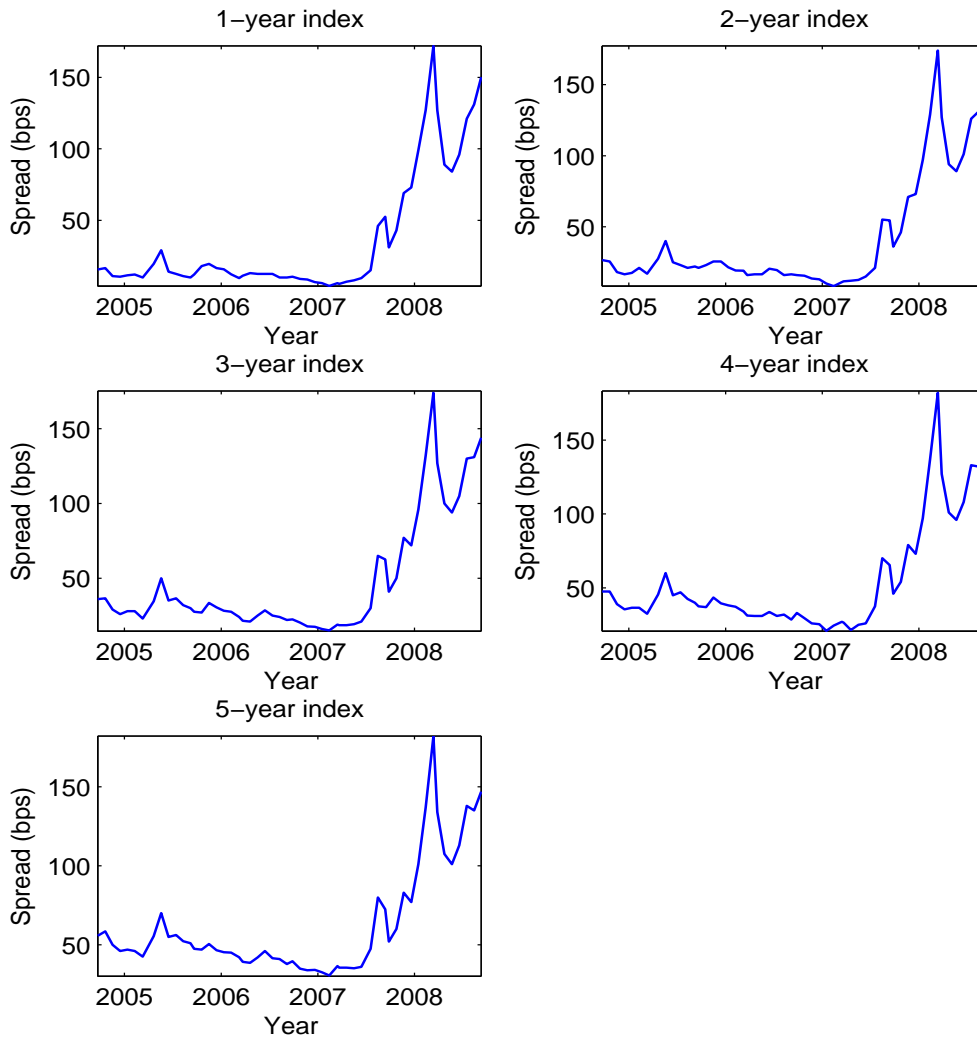


Figure 4: Historical time series of spreads for the one-year, two-year, three-year, four-year and five-year CDX indices. Our benchmark model is calibrated to perfectly match these time series.

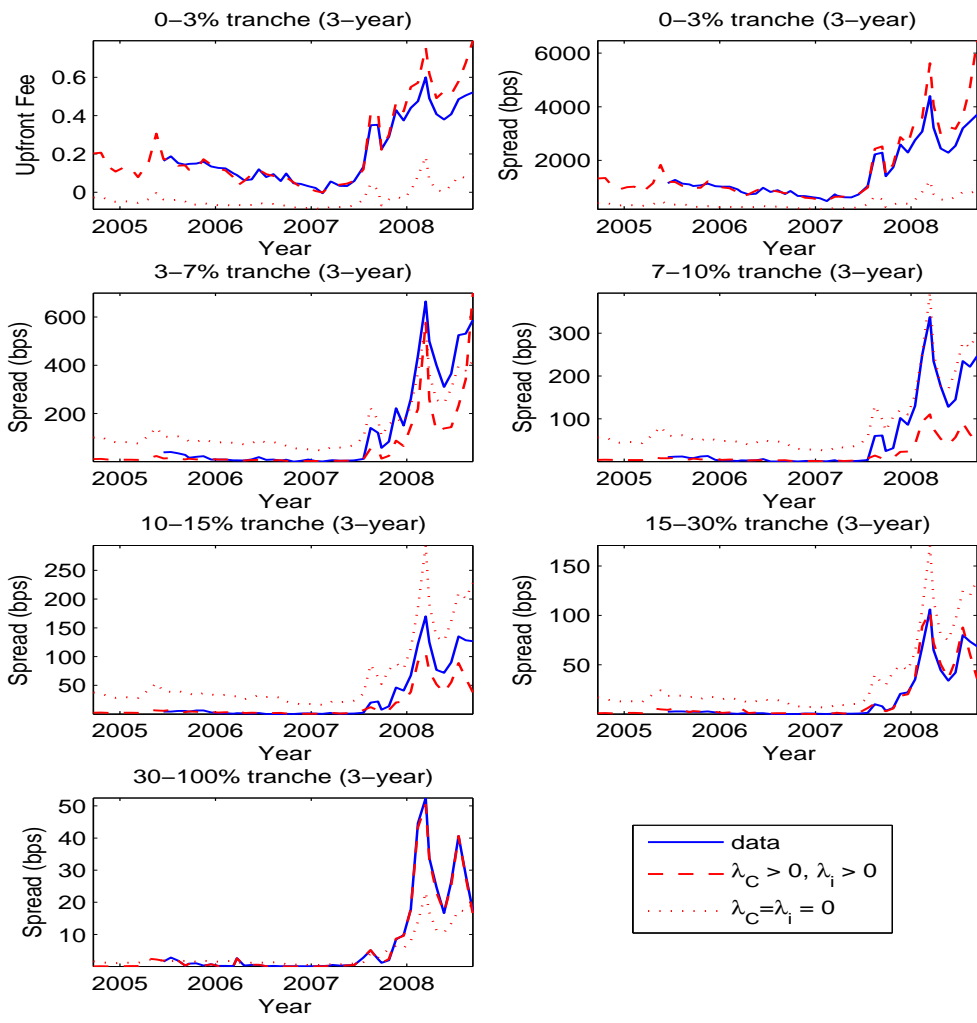


Figure 5: Predicted time series of spreads for the 0-3% up-front premium, 0-3% running premium, 3-7%, 7-10%, 10-15%, 15-30% three-year CDX tranches for various model specifications. ‘data’ is the historical data. $\lambda_C > 0, \lambda_i > 0$ denotes our benchmark model with idiosyncratic jumps fitted to the 1 to 5 year term structure of CDX, and catastrophe jump intensity fitted to the super senior tranche. $\lambda_C = 0, \lambda_i = 0$ has neither catastrophe nor idiosyncratic jump risk.

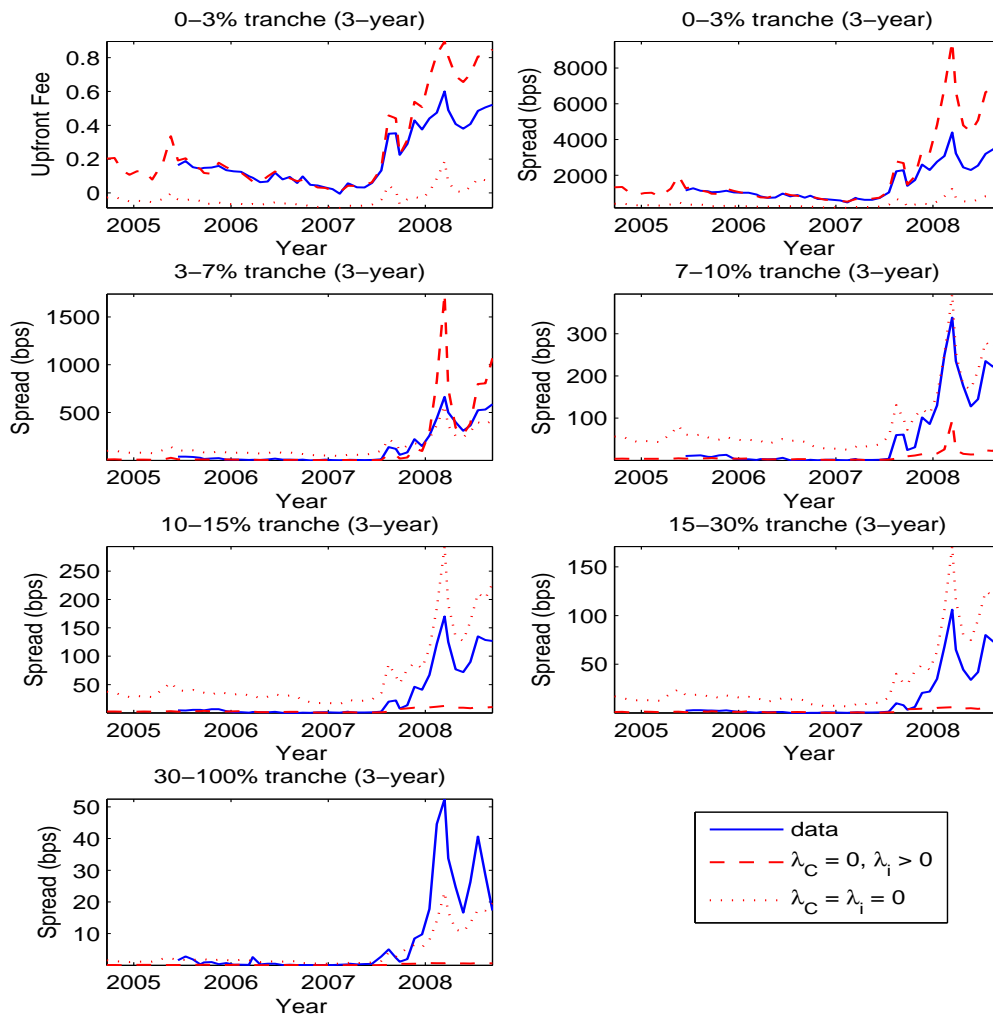


Figure 6: Predicted time series of spreads for the 0-3% up-front premium, 0-3% running premium, 3-7%, 7-10%, 10-15%, 15-30% three-year CDX tranches for various model specifications. ‘data’ is the historical data. $\lambda_C = 0, \lambda_i > 0$ denotes our model with idiosyncratic jumps fitted to the 1 to 5 year term structure of CDX, but where we do not allow for catastrophe jumps (i.e., systematic jumps are extracted solely from index options). $\lambda_C = 0, \lambda_i = 0$ has neither catastrophe nor idiosyncratic jump risk.

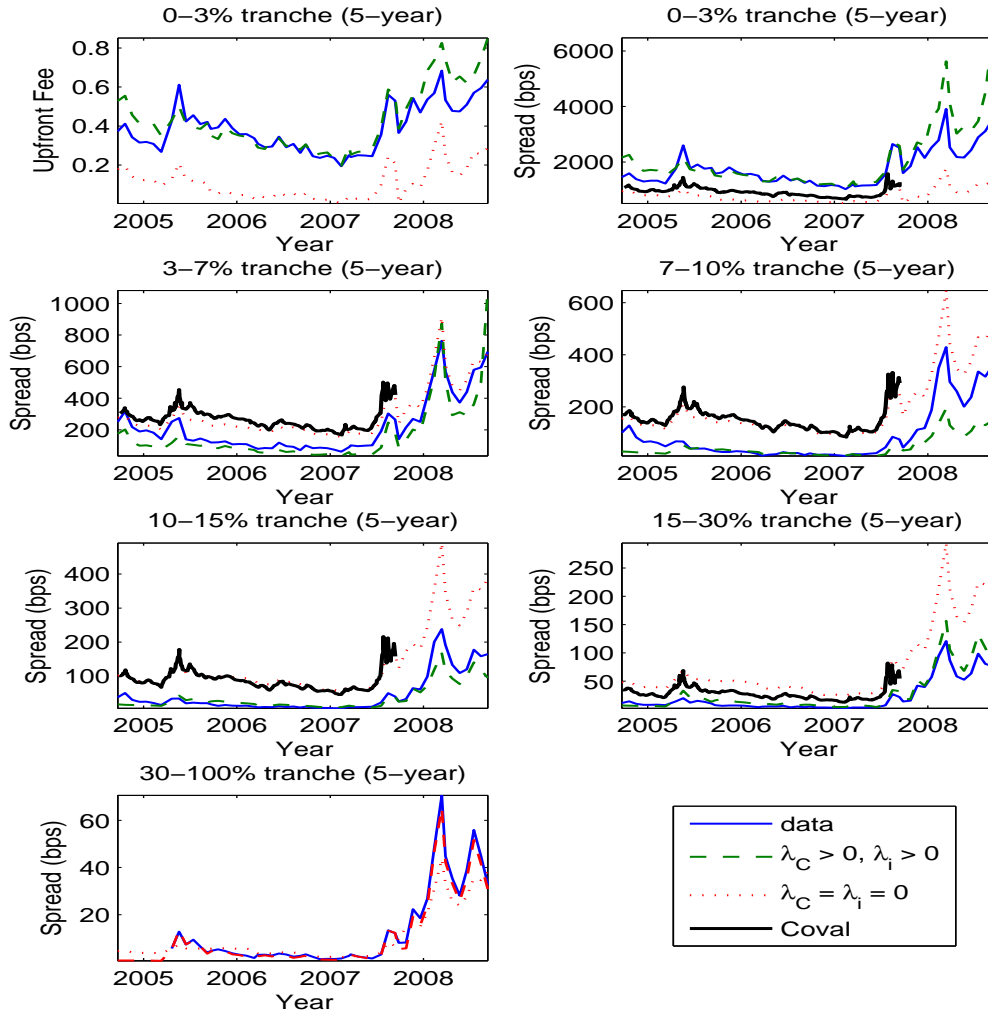


Figure 7: Predicted time series of spreads for the 0-3% up-front premium, 0-3% running premium, 3-7%, 7-10%, 10-15%, 15-30% five-year CDX tranches for various model specifications. ‘data’ is the historical data. $\lambda_C > 0, \lambda_i > 0$ denotes our benchmark model with idiosyncratic jumps fitted to the 1 to 5 year term structure of CDX, and catastrophe jump intensity fitted to the super senior tranche. $\lambda_C = 0, \lambda_i = 0$ has neither catastrophe nor idiosyncratic jump risk. When available we also plot the tranche spreads predicted by CJS (denoted by ‘coval’).

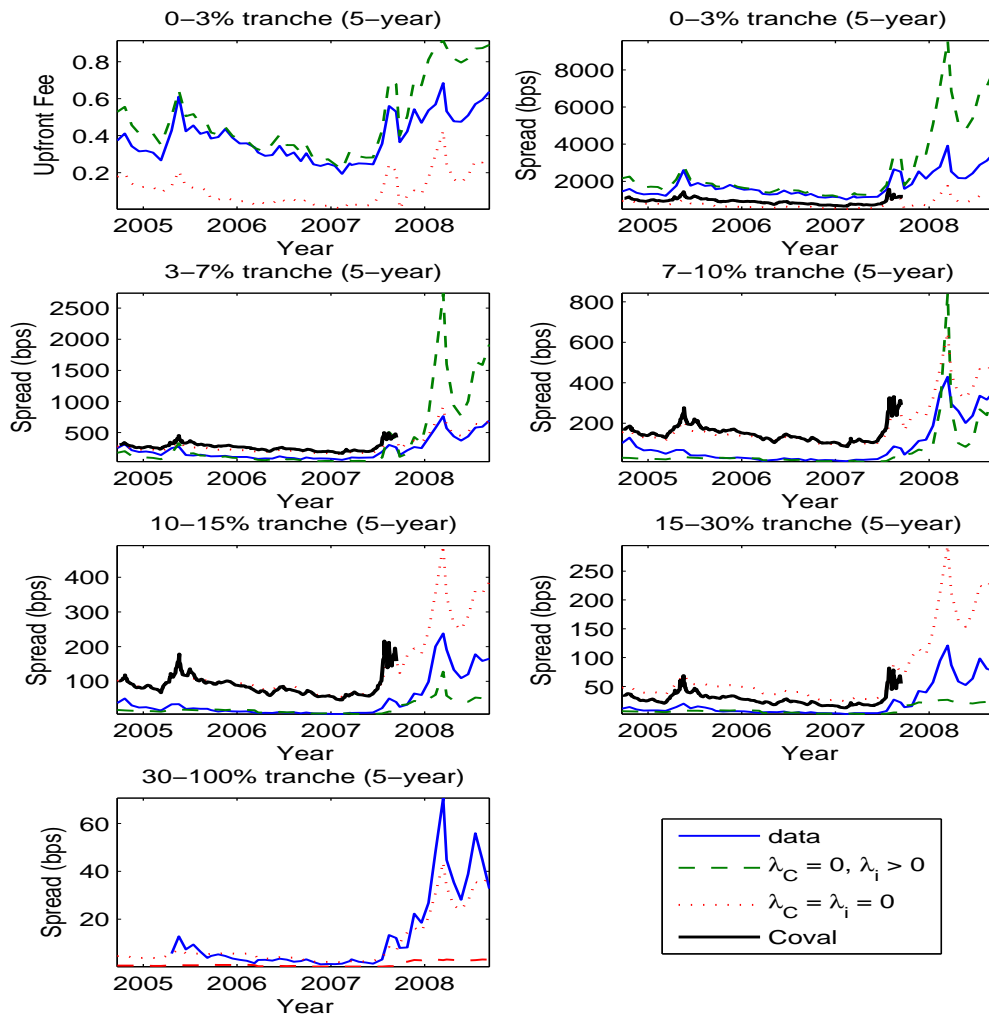


Figure 8: Predicted time series of spreads for the 0-3% up-front premium, 0-3% running premium, 3-7%, 7-10%, 10-15%, 15-30% five-year CDX tranches for various model specifications. ‘data’ is the historical data. $\lambda_C = 0, \lambda_i > 0$ denotes our model with idiosyncratic jumps fitted to the 1 to 5 year term structure of CDX, but where we do not allow for catastrophe jumps (i.e., systematic jumps are extracted solely from index options). $\lambda_C = 0, \lambda_i = 0$ has neither catastrophe nor idiosyncratic jump risk. When available we also plot the tranche spreads predicted by CJS (denoted by ‘coval’).

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] G. Andrade and S. N. Kaplan. How costly is financial (not economic) distress? evidence from highly leveraged transactions that became distressed. *Journal of Finance*, v53 n5:1443–1494, 1998.
- [4] D. Backus, M. Chernov, and I. Martin. Disasters implied by equity index options. Forthcoming *Journal of Finance*, 2011.
- [5] D. Backus, S. Foresi, and L. Wu. Accounting for biases in black-scholes. *working paper*, 2004.
- [6] J. Bai and P. Collin-Dufresne. The determinants of the cds-bond basis during the financial crisis of 2007-2009. Working paper, 2011.
- [7] 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.
- [8] R. J. Barro. Rare disasters and asset markets in the twentieth century. *Quarter Journal of Economics*, 121(3):823–866, 2006.
- [9] D. S. Bates. Post-'87 crash fears in the s&p 500 futures option market. *Journal of Econometrics*, 94:181–238, 2000.
- [10] F. Black and J. C. Cox. Valuing corporate securities: some effects of bond indentures provisions. *Journal of Finance*, 31:351–367, 1976.
- [11] 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.
- [12] D. T. Breeden and Litzenberger. Prices of stochastic contingent claims implicit in option prices. *Journal of Business*, 51:621–651, 1978.

- [13] M. Broadie, M. Chernov, and M. Johannes. Model specification and risk premia: Evidence from futures options. *Journal of Finance*, 62:1453–1490, 2007.
- [14] P. Carr and D. Madan. Option valuation using the fast fourier transform. *Journal of Computational Finance*, 2:61–73, 1999.
- [15] P. Carr and L. Wu. The finite moment log stable process and option pricing. *Journal of Finance*, 58:753–777, 2003.
- [16] P. Christoffersen, S. Heston, and K. Jacobs. The shape and term structure of the index option smirk: Why multifactor stochastic volatility models work so well. *Management Science*, 55(12):1914–1932, 2009.
- [17] P. Collin-Dufresne. A short introduction to correlation markets. *Journal of Financial Econometrics*, 7(1):12–29, 2009.
- [18] P. Collin-Dufresne and R. S. Goldstein. Do credit spreads reflect stationary leverage ratios? *Journal of Finance*, 56:1929–1957, 2001.
- [19] P. Collin-Dufresne, R. S. Goldstein, and J. Helwege. Is credit event risk priced? modeling contagion via the updating of beliefs. *working paper*, 2003.
- [20] J. D. Coval, J. W. Jurek, and E. Stafford. Economic catastrophe bonds. *American Economic Review*, 99(3):628–666, 2009.
- [21] 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.
- [22] S. R. Das, D. Duffie, N. Kapadia, and L. Saita. Common failings: How corporate defaults are correlated. *Journal of Finance*, pages 93–117, 2007.
- [23] S. R. Das, L. Freed, G. Geng, and N. Kapadia. Correlated default risk. *Journal of Fixed Income*, pages 7–32, 2006.
- [24] S. R. Das and R. Sundaram. Of smiles and smirks: A term structure perspective. *Journal of Financial and Quantitative Analysis*, 34:211–240, 1999.
- [25] S. A. Davydenko. What triggers default? a study of the default boundary. Working paper, 2010.

- [26] S. A. Davydenko, I. A. Strebulaev, and X. Zhao. A market-based study of the costs of default. Working paper, 2011.
- [27] E. Derman and I. Kani. Riding on a smile. *RISK*, 7:32–39, 1994.
- [28] D. Duffie. Credit swap valuation. *Financial Analysts Journal*, 55n1:73–87, 1999.
- [29] D. Duffie and N. Garleanu. Risk and valuation of collateralized debt obligations. *Financial Analysts Journal*, 57n1:41–59, 2001.
- [30] D. Duffie, J. Pan, and K. Singleton. Transform analysis and option pricing for affine jump-diffusions. *Econometrica*, 68:1343–1376, 2000.
- [31] D. Duffie, L. Saita, and K. Wang. Multi-period corporate default prediction with stochastic covariates. *Journal of Financial Economics*, 83:635–665, 2007.
- [32] B. Dumas, J. Fleming, and R. Whaley. Implied volatility functions: Empirical tests. *Journal of Finance*, 53:2059–2106, 1998.
- [33] B. Dupire. Pricing with a smile. *RISK*, 7:18–20, 1994.
- [34] A. Eckner. Computational techniques for basic affine models of portfolio credit risk. *Journal of Computational Finance*, 13:63–97, 2009.
- [35] Y. Eom, J. Helwege, and J.-Z. Huang. Structural models of corporate bond pricing. *RFS*, 17:499–544, 2004.
- [36] B. Eraker. Do stock prices and volatility jump? reconciling evidence from spot and option prices. *Journal of Finance*, 59:1367–1403, 2004.
- [37] B. Eraker, M. S. Johannes, and N. Polson. The impact of jumps in volatility and returns. *Journal of Finance*, 58:1269–1300, 2003.
- [38] K. Giesecke and L. R. Goldberg. A top down approach to multi-name credit. *working paper*, 2005.
- [39] R. Goldstein, N. Ju, and H. Leland. An ebit-based model of dynamic capital structure. Forthcoming, *Journal of Business*, 2001.
- [40] M. Harrison and D. Kreps. Martingales and multiperiod securities markets. *Journal of Economic Theory*, 20:381–408, 1979.

- [41] S. L. Heston. A closed form solution for options with stochastic volatility. *Review of financial studies*, 6:327–343, 1993.
- [42] 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.
- [43] J. Hull and A. White. The risk of tranches created from mortgages. *Financial Analysts Journal*, 66(5):54–67, 2010.
- [44] 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.
- [45] P. Jorion and G. Zhang. Credit contagion and counterparty risk. *Journal of Finance*, 64:2053–2087, 2009.
- [46] H. Leland. Predictions of default probabilities in structural models of debt. *Journal of Investment Management*, 2(2), 2004.
- [47] D. X. Li. On default correlation: A copula function approach. *Journal of Fixed Income*, v9n4:43–54, 2000.
- [48] 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.
- [49] F. A. Longstaff and A. Rajan. An empirical analysis of the pricing of collateralized debt obligations. *Journal of Finance*, v63n2:529–563, 2008.
- [50] R. C. Merton. On the pricing of corporate debt: The risk structure of interest rates. *Journal of Finance*, 29:449–470, 1974.
- [51] A. Mortensen. Semi-analytical valuation of basket credit derivatives in intensity-based models. *Journal of Derivatives*, v13n4:8–26, 2006.
- [52] C. Nelson and A. Siegel. Parsimonious modeling of yield curves. *Journal of Business*, 60:473–489, 1987.
- [53] J. Pan. The jump-risk premia implicit in options: evidence from an integrated time-series study. *Journal of Financial Economics*, 63:3–50, 2002.

- [54] T. A. Rietz. The equity risk premium a solution. *Journal of Monetary Economics*, 22:117–131, 1988.
- [55] M. Rubinstein. Implied binomial trees. *Journal of Finance*, 49:771–818, 1994.
- [56] S. M. Schaefer and I. A. Strebulaev. Structural models of credit risk are useful: Evidence from hedge ratios on corporate bonds. *Journal of Financial Economics*, 90:1–19, 2008.
- [57] O. Vasicek. An equilibrium characterization of the term structure. *Journal of Financial Economics*, 5:177–188, 1977.
- [58] O. Vasicek. Probability of loss on a loan portfolio. *Moody's-KMV Research document*, 1987.
- [59] C. Zhou. An analysis of default correlation and multiple defaults. *The Review of Financial Studies*, v14(2):555–576, 2001.

7 Appendices

7.1 Numerical solution of the long-term option model

We add two ingredients to the standard SVCJ option model. First, a second stochastic market variance to calibrate the term structure of option implied volatilities.²³ Second, a catastrophe jump is introduced to match the super-senior tranche (30-100%) of the CDX. The log market dynamics ($m_t = \log M_t$) are

$$dm_t = \left(r - \delta - \bar{\mu}_y \lambda^Q - (e^{y_C} - 1) \lambda_C^Q - \frac{1}{2} V_t - \frac{1}{2} \theta_t \right) dt + \sqrt{V_t} dw_1^Q + \sqrt{\theta_t} dw_2^Q + y dq + y_C dq_C. \quad (17)$$

$$dV_t = \kappa_V (\bar{V} - V_t) dt + \sigma_V \sqrt{V_t} (\rho_1 dw_1^Q + \sqrt{1 - \rho_1^2} dw_3^Q) + y_V dq \quad (18)$$

$$d\theta_t = \kappa_\theta (\bar{\theta} - \theta_t) dt + \sigma_\theta \sqrt{\theta_t} (\rho_2 dw_2^Q + \sqrt{1 - \rho_2^2} dw_4^Q) + y_\theta dq. \quad (19)$$

Since the dynamics are affine, the moment generating function of the log market is

$$\phi_T(u) = \mathbb{E}_0^Q [e^{um_T}] = e^{A(T) + um_0 + B(T)V_0 + C(T)\theta_0},$$

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

$$\begin{aligned} A'(t) &= (r - \delta - \bar{\mu}_y \lambda^Q - (e^{y_C} - 1) \lambda_C^Q) u + \kappa_V \bar{V} B(t) + \kappa_\theta \bar{\theta} C(t) \\ &= + \lambda^Q \mathbb{E}^Q \left[e^{uy + B(t)y_V + C(t)y_\theta} - 1 \right] + \lambda_C^Q (e^{y_C} - 1) \end{aligned} \quad (20)$$

$$B'(t) = -\frac{1}{2}u + \frac{1}{2}u^2 - \kappa_V B(t) + \frac{1}{2}\sigma_V^2 B(t)^2 + \rho_1 \sigma_V u B(t) \quad (21)$$

$$C'(t) = -\frac{1}{2}u + \frac{1}{2}u^2 - \kappa_\theta C(t) + \frac{1}{2}\sigma_\theta^2 C(t)^2 + \rho_2 \sigma_\theta u C(t) \quad (22)$$

where

$$\mathbb{E}^Q \left[e^{uy + B(t)y_V + C(t)y_\theta} - 1 \right] = \frac{e^{u\mu_y + \frac{1}{2}u^2\sigma_y^2}}{(1 - \mu_V B(t))(1 - \mu_\theta C(t))} - 1. \quad (23)$$

The boundary conditions are

$$A(0) = B(0) = C(0) = 0. \quad (24)$$

We apply the FFT to the characteristic function $\phi_T(iv)$ to get the risk-neutral market distribution at horizon 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).

²³In the robustness check of stochastic interest rates, we have an additional state variable of short-term rate.

7.2 Calibration procedure for Market model

The details of the calibration of market dynamics are described below. We bootstrap daily zero-coupon curve from swap rates of maturities 1, 2, 3, 4, 5, 6-year using the extended Nelson-Siegel approach (Nelson and Siegel (1987)). Market dividend yields are from OptionMetrics. Other parameters are chosen to fit five-year option implied volatilities with moneyness from 0.5 to 1.5 with 0.1 increments and one-year at-the-money option. Option prices are computed using the Fast Fourier Transformation. When fitting these options, we set $\rho_1 = -0.48$ following BCJ²⁴, and optimize other parameters (Θ) of the option model to minimize the overall relative Root Mean Square Errors (RMSE) of implied volatilities for every series:

$$\Theta^* = \arg \min_{\Theta} \left\{ \sum_t RMSE_t(\Theta; V_t^*, \theta_t^*)^2 \right\}. \quad (25)$$

Here, the overall relative RMSE is computed as the sum of the minimized relative RMSE of the implied volatilities (σ_n) of one-year at-the-money option and eleven five-year options for each date (t) over state variables (V_t and θ_t) as

$$RMSE_t(\Theta; V_t^*, \theta_t^*) = \sqrt{\min_{V_t, \theta_t} \sum_n w_n \left(\frac{\sigma_{n,model}(\Theta; V_t, \theta_t)}{\sigma_{n,data}} - 1 \right)^2}. \quad (26)$$

In this formula, (V_t^*, θ_t^*) are the optimized state variables, and w_n is the weight on option n . We put more weight on fitting the one-year at-the-money option than five-year options to perfectly match the level of short-term market volatility.

For the model with crash-risk intensity calibrated to super senior tranches, we adopt an iterative procedure starting from the parameters obtained without crash risk. Because of the heavy computational costs of pricing the super senior tranche using Monte Carlo simulation, we use a relatively simple approach to calibrate λ_C^Q . We assume that the catastrophic jump intensity is piece-wise linear: a constant equal to $\lambda_{C,3}^Q$ for the first three years, and equal to $\lambda_{C,5}^Q$ henceforth. We first use the parameters obtained with $\lambda_C^Q = 0$ and we simulate the model to compute super senior tranche spreads (S_L, S_H) for two given catastrophe jump intensities ($\lambda_{C,L}^Q, \lambda_{C,H}^Q$). Then, for each super-senior spread observed in the data (S), we compute the $\lambda_{C,i}^Q$ parameters using linear interpolations. Specifically, we find $\lambda_{C,3}^Q$ to match the three-year super senior tranche spread first and then $\lambda_{C,5}^Q$ to match the five-year super senior tranche spread (given $\lambda_{C,3}^Q$). Lastly, we recalibrate option parameters and iterate if necessary. We find that the

²⁴BCJ advocate using some time-series information when theory restricts risk-neutral and physical measure parameters to be the same and the latter are fairly well estimated from data. We found this parameter to be difficult to pin down from the cross-section alone, and therefore use BCJ's estimate.

super senior tranche spreads are fairly insensitive to changes in the remaining parameters and state variables. Therefore, this approach works fairly well with a single set of $\lambda_{C,L}^Q$ and $\lambda_{C,H}^Q$ even in time-series calibrations. In fact, we find that during the pre-crisis period, we do not even need to iterate the calibration procedure to achieve a good fit. Therefore we report only one set of parameters for the non-crash related parameters. During the crisis period however, we report two sets of parameters, as there the interaction between crash parameters and other parameters is more significant. In Figure 9 below we present the time series of $RMSE_t$ as well as the implied volatility state variable dynamics for V_t and θ_t . Further, as a check that our calibration is sensible, we report on Figure 10 the fitted instantaneous volatility ($\sqrt{V_t + \theta_t}$) and compare it with the VIX index. While these should not be exactly the same, we can expect a high correlation between them, and it is therefore comforting to see that they track each other fairly closely. Lastly, in Figure 11, we report the time series of crash intensities. It is clear that before the crisis the crash-intensity is close to zero. There is a small increase around the May 2005 period corresponding to the downgrade of Ford and GM. But mostly the intensity is insignificantly different from zero. However mid-2007 we see a tremendous increase in the crash intensity to 100 bps in 2008. We emphasize that the five-year line on the graph really corresponds to the three into five forward intensity, given the way we parameterize the intensity step function. The graph therefore also shows that the perceived forward crash intensities are quite a bit higher both during the May 2005 event and the crisis, indicating that this affects more the five-year super senior than the three-year. At its maximum, then the super-senior implied catastrophe occurs once every hundred years, with a magnitude that far exceeds even the great depression (since super-senior insurance pays off only if more than 37.5 percent of the investment grade portfolio defaults, given our twenty percent recovery assumption).

7.3 Estimation of firm level statistics

Here we explain how we construct estimate the quantities reported in Table 2.

The risk-free rate and S&P 500 index dividend yield are time-series averages for each six month period corresponding to the series. Table 2 reports these estimates over series in our sample.

The leverage ratio is defined as book debt divided by the sum of book debt and market equity, where book debt is computed by the sum of short-term debt and long-term debt from quarterly Compustat. If any of these numbers are missing, we use their corresponding items in annual data. Table 2 reports the leverage ratio averages over firms and time-series of a five-year moving window that ends at the initial date of every series.

We estimate asset betas and idiosyncratic volatilities using return data of every collateral firms. Similar to the method of estimating asset volatility in Schaefer and Strebulaev (2008), we first compute asset returns using average of equity returns and debt returns weighted by leverage ratios and then we use CAPM regression to estimate asset betas and idiosyncratic volatilities. Because all collateral firms are investment-graded, we use the returns on LQD (an Exchange-Traded Fund of investment-graded corporate bonds) to proxy for debt returns. S&P 500 index excess returns are used as the market factor. Within a five-year moving window that ends at the initial date of every series, we regress the time-series of asset excess returns of each firms on the market factor to estimate asset betas. Asset idiosyncratic volatility is estimated as the standard deviation of the regression residuals.²⁵ Averages of asset betas and idiosyncratic volatilities over firms for every series are reported in Table 2.

The aggregate asset payout ratio of collateral firms is estimated by taking average of aggregate equity dividend yield and average corporate bond yield weighted by the average leverage ratio of every series. Aggregate equity dividend yield is defined as S&P 500 index dividend yield. And average corporate bond yield is calculated by summing the risk-free rate and average investment-graded CDS spread, approximately the five-year CDX index spread averaging over series.

7.4 Calibration of idiosyncratic intensity parameters

Here we explain how we proceed to calibrate the idiosyncratic intensity step function to the one to five year term structure of CDX credit spreads. We define the state variable vector of our model by $X_t = \{x_{1,t}, x_{2,t}, \dots, x_{9,t}\}$, where $x_{i,t}$ represents idiosyncratic jump intensities (λ_j^Q with $j = 1-, 2-, 3-, 4-, 5\text{-year}$), three-year catastrophe jump intensity ($\lambda_{C,3}^Q$), the five-year catastrophe jump intensity ($\lambda_{C,5}^Q$), and market variances (V_t and θ_t). First, we set $\lambda_{C,3}^Q$, $\lambda_{C,5}^Q$, V_t , and θ_t to their time-series average values and calibrate $\bar{\lambda}_j^Q$ to the average CDX index using linear interpolation similarly to the calibration of super senior tranches. Second, we alter each state variable $x_{i,t}$ to a region ($x_{i,L}$ and $x_{i,H}$) and simulate the model to price CDX indices ($S_{idx,j}$) with all maturities ($j = 1-, 2-, 3-, 4-, 5\text{-year}$). Third, we compute the CDX indices sensitivity matrix (M) to state variables by regressing different sets of $S_{idx,j}$ with $j = 1-, 2-, 3-, 4-, 5\text{-year}$ on their corresponding state variables (X_t). Fourth, we invert the part of M , which corresponds to idiosyncratic jump intensities, to back out λ_j^Q from CDX index data with other

²⁵This estimate tends to underestimate asset idiosyncratic volatility. Real idiosyncratic volatility should be higher, which reduces senior tranche spreads even more. LQD is an aggregate index of investment-graded bond returns. Idiosyncratic bond volatility is diversified. Therefore, when we use it to proxy the return on debt of the collateral firms, we underestimate the asset idiosyncratic volatility.

state variables calibrated as described in previous sections. To accelerate the process, we only simulate one firm to price the CDX index spreads rather than 125 firms because CDX indices only depends on marginal distribution of defaults.

The estimated time series of idiosyncratic jump intensities are shown in figures 12 and 13 (which respectively correspond to the calibration with and without Market catastrophe risk). As can be seen when comparing to Figure 4 it is clear that the one and two year intensities directly capture the credit spread risk embedded in short term CDX index spreads (indeed, the approximation that the spread is equal to the risk-neutral intensity times loss given default works almost perfectly). At the long end the relation is not as tight, as one would expect because of the possibility of diffusion risk triggering default. When comparing the calibration with and without catastrophic risk, one observes that the one-year intensity in the case without catastrophic risk is basically the sum of the idiosyncratic intensity estimated when there is catastrophic risk and the catastrophic intensity, as one would expect, since in the latter case, there are two sources of jump risk that can lead to early default. The presence of sizable catastrophic risk reduces the size of the idiosyncratic jump intensity required to match the large short term credit spreads. We also see from the term structure of intensities, that during the pre-crisis period the intensities are fairly flat or increasing, whereas during the crisis period the intensities are flat or decreasing as a function of maturity, reflecting different slopes in the term structure of index spreads.

Figures 14 and 15 show the estimated total idiosyncratic risk (jump and diffusion), versus the pure diffusion idiosyncratic risk (based on Table 2). We note that pre-crisis, the fact that we calibrate diffusion idiosyncratic risk to match the entire idiosyncratic risk estimated from past time series data, leads to only a small difference in our model idiosyncratic risk relative to the time series estimates, since the estimated idiosyncratic jump intensities are rather small. However, during the crisis since credit spreads widen tremendously across all maturities and much more so than the systematic jump risk implied from the options data, estimated idiosyncratic jump intensities are quite large (of the order of 200bps). In turn, this leads to a much larger estimate of idiosyncratic risk during the crisis (about 50%) larger than based on our pre-crisis time series estimates.

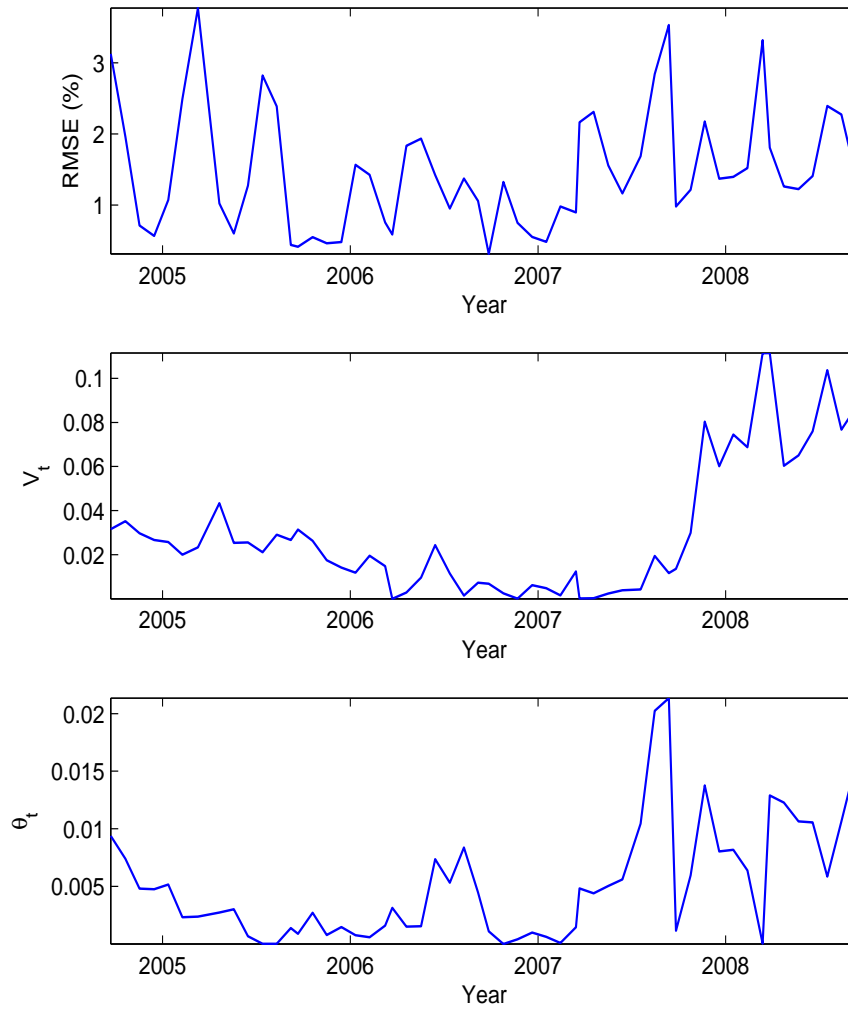


Figure 9: Estimated volatility state variables inverted from ATM one-year and eleven ITM, ATM and OTM five-year option prices based on our option pricing model and using the parameters presented in Table 1. We also plot the RMSE for the panel data of options we fit every day.

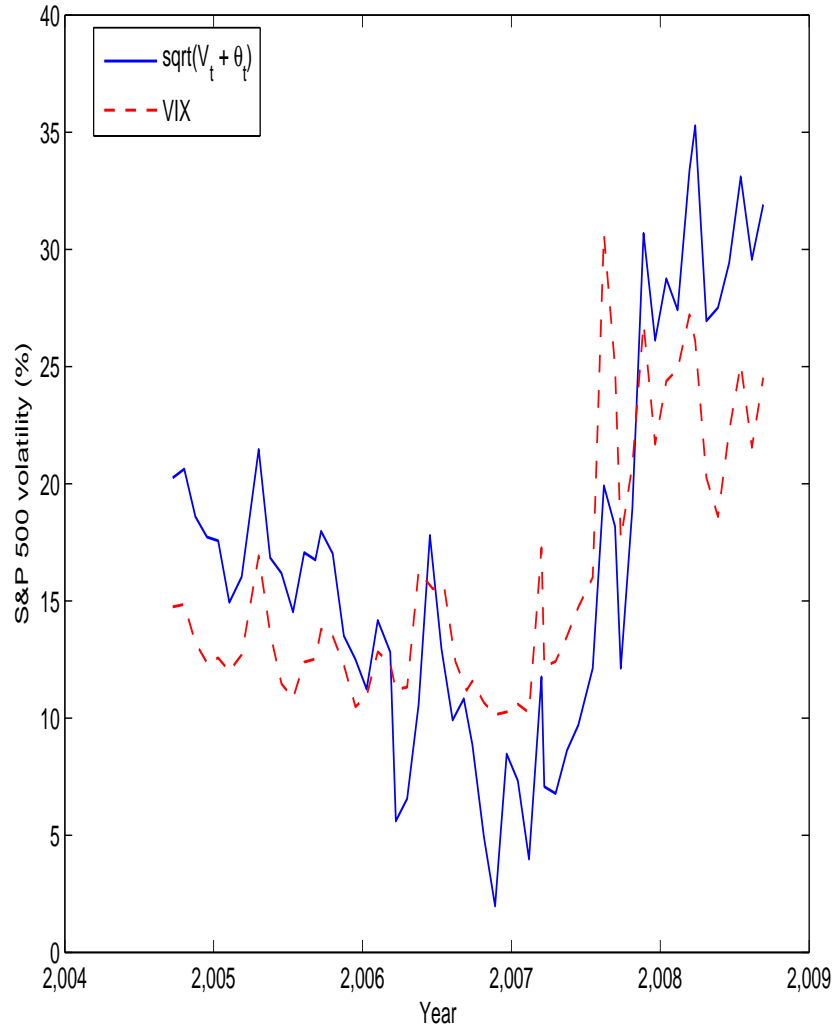


Figure 10: Estimated instantaneous volatility $\sqrt{V_t + \theta_t}$ of the market return, where the volatility state variables V_t and θ_t are inverted from ATM one-year and eleven ITM and OTM five-year maturity option prices based on our option pricing model and using the parameters presented in Table 1. We also plot the VIX along side our instantaneous volatility to give a sense of the reasonableness of the fitted latent volatility state variables.

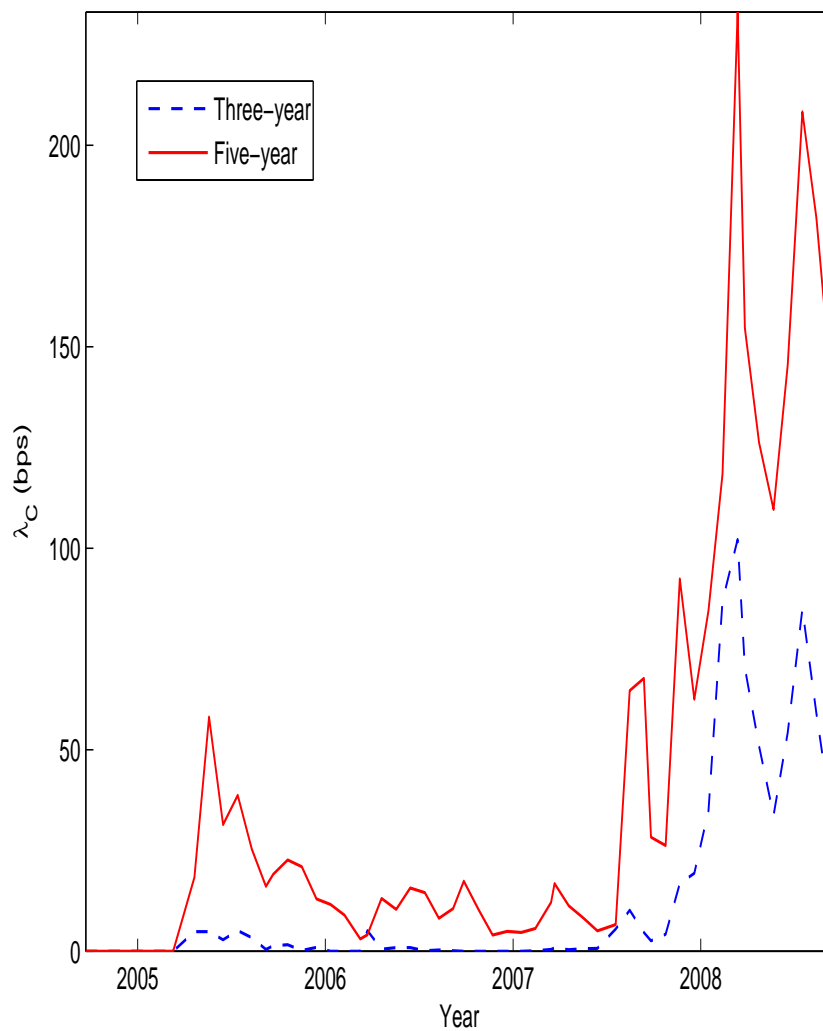


Figure 11: Estimated crash risk intensity. As explained in section 3.2 we calibrate λ_C^Q to match perfectly super senior tranches with three and five year maturity as well as S&P500 index options.

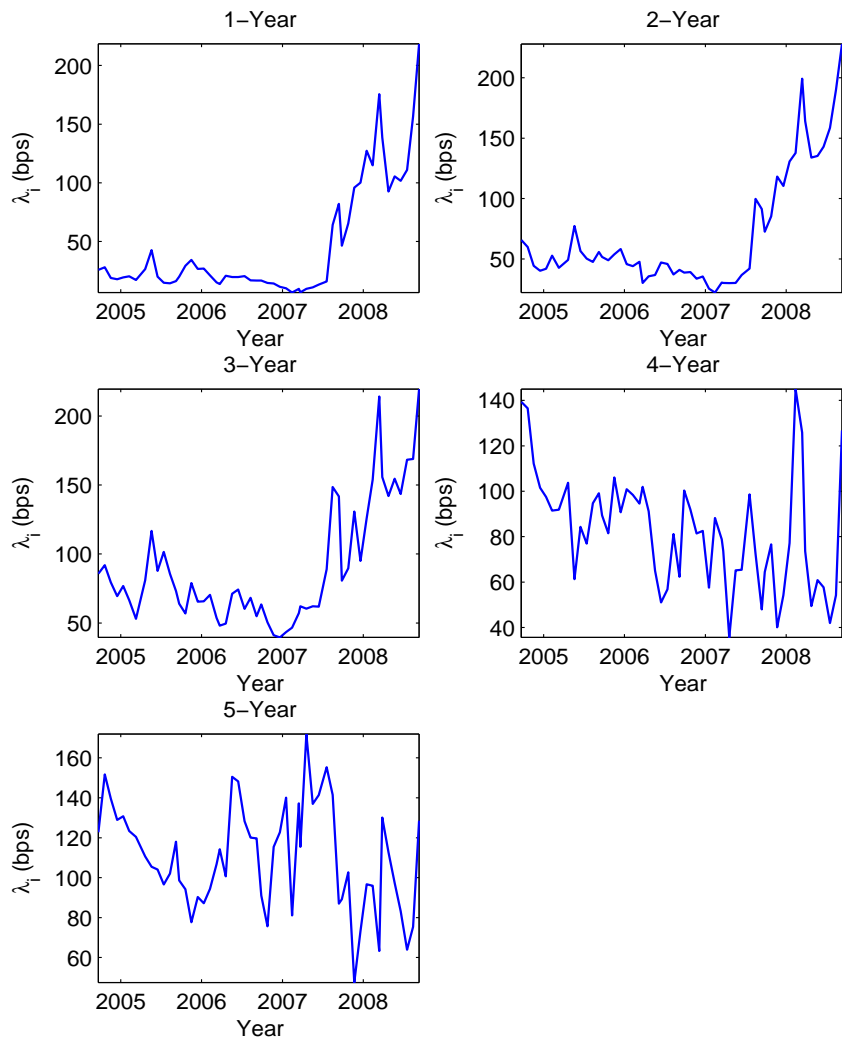


Figure 12: Estimated idiosyncratic jump risk intensity when using the benchmark model with catastrophe risk fitted to super-senior tranches ($\lambda_C > 0$)

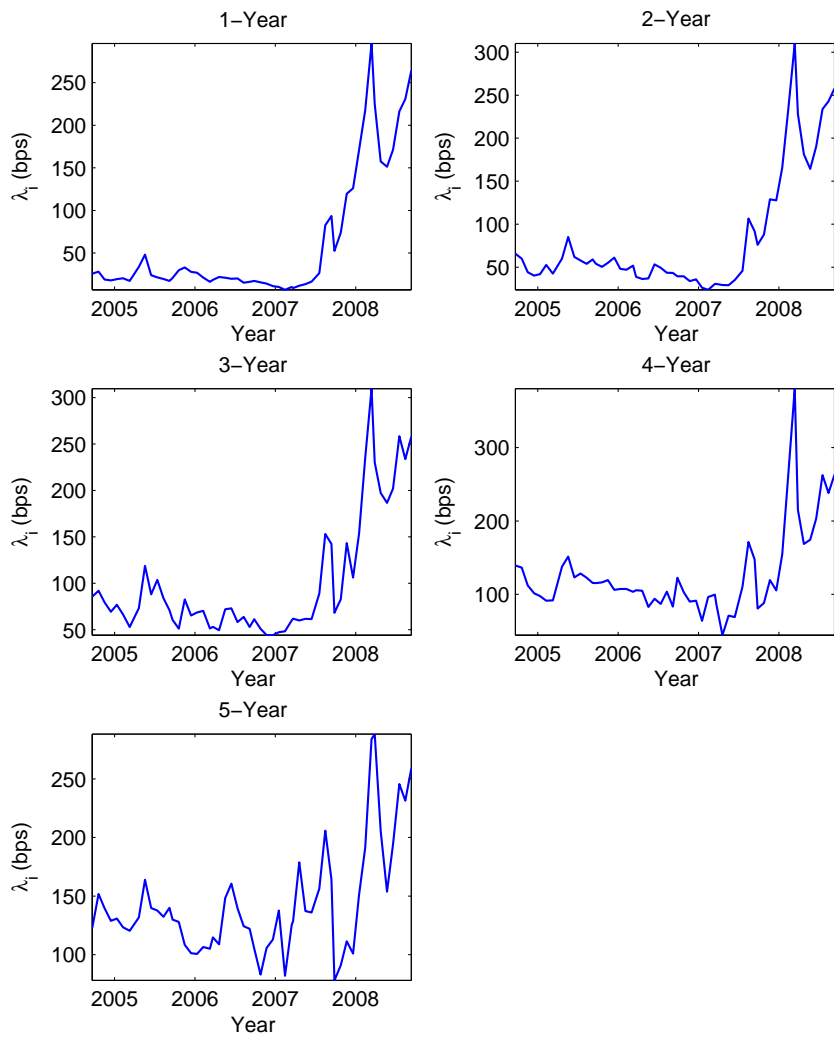


Figure 13: Estimated idiosyncratic jump risk intensity when using the model without catastrophe risk ($\lambda_C = 0$ and all systematic jump risk is extracted from index options alone)

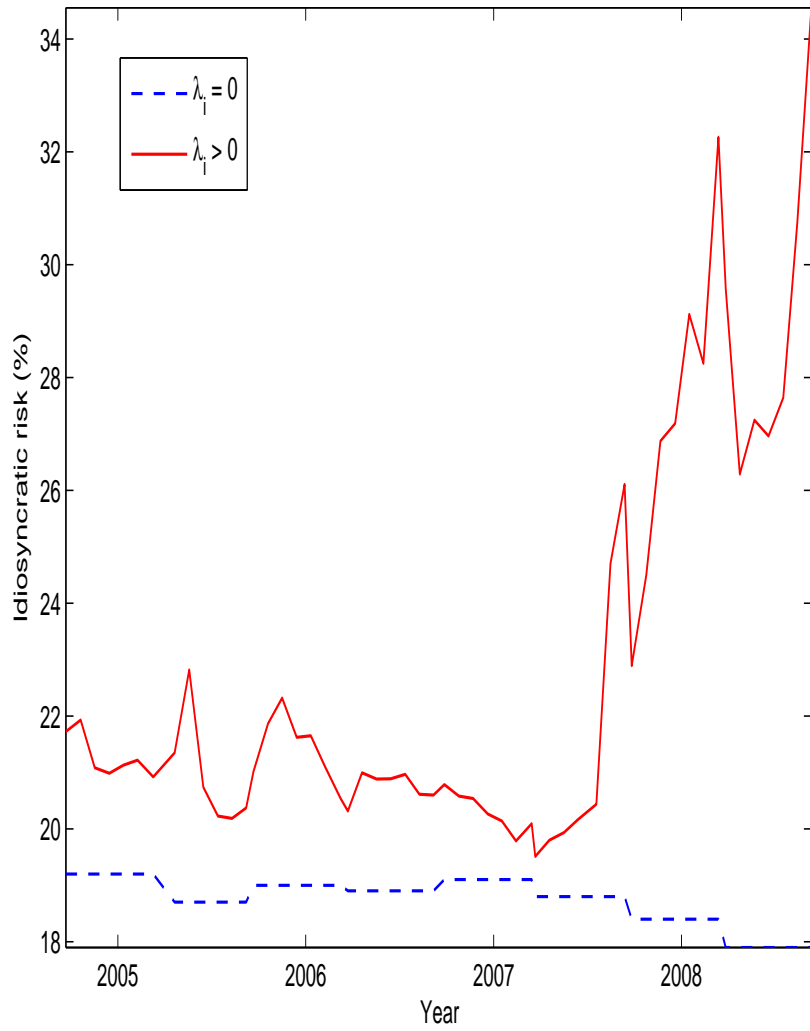


Figure 14: Impact of idiosyncratic jumps on estimated idiosyncratic risk in the benchmark model with catastrophe risk fitted to super-senior tranches ($\lambda_C > 0$)

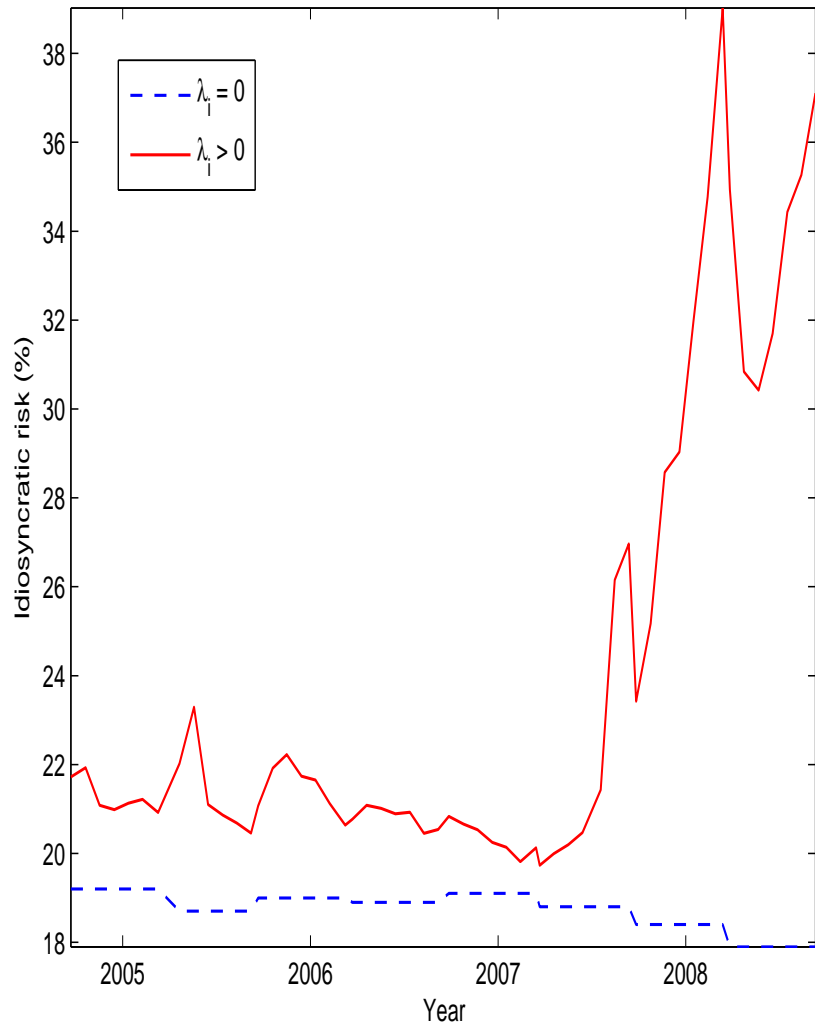


Figure 15: Impact of idiosyncratic jumps on estimated idiosyncratic risk in the model without catastrophe risk ($\lambda_C = 0$ and all systematic jump risk is extracted from index options alone)