Do rare events explain CDX tranche spreads?*

Sang Byung Seo
University of Houston

Jessica A. Wachter
University of Pennsylvania
and NBER

September 16, 2016

Abstract
We investigate whether a model with a time-varying probability of economic disaster can explain the pricing of collateralized debt obligations, both prior to and during the 2008-2009 financial crisis. Namely, we examine the pricing of tranches on the CDX, an index of credit default swaps on large investment-grade firms. CDX senior tranches are essentially deep out-of-the money put options because they do not incur losses until a large fraction of previously stable firms default. As such, these products clearly reflect the market's assessment of rare-event risk. We find that the model can simultaneously explain prices on CDX senior tranches and on equity index options at parameter values that are consistent with the equity premium and with aggregate stock market volatility. Our results demonstrate the importance of beliefs about rare disasters for asset prices, even during periods of relative economic stability.

*Seo: sseo@bauer.uh.edu; Wachter: jwachter@wharton.upenn.edu. We thank Pierre Collin-Dufresne, Hitesh Doshi, Kris Jacobs, Mete Kilic, Nick Roussanov, Ivan Shaliastovich, Pietro Veronesi, Amir Yaron and seminar participants at AFA annual meetings, University of Houston, and the Wharton School for helpful comments.

1 Introduction

The period from 2005 to September of 2008 witnessed a more than 100-fold increase in the cost of insuring against economic catastrophe. This cost can be seen from the pricing of derivative contracts written on the CDX, an index of credit default swaps on investment-grade firms. During the 2005–2009 period, tranches on the CDX were actively traded: investors could purchase and sell insurance that would pay off only if a certain fraction of firms represented by the CDX went into default. The most senior tranches were structured to pay off only if corporate defaults became extremely widespread, more so than during the Great Depression.

While the costs to insure the senior tranches on the CDX were close to zero through most of 2006 and 2007, fluctuations began to appear in late 2007, culminating in sharply rising prices in the summer and fall of 2008. Ex post, of course, such insurance did not pay off. In fact only a very small number of firms represented by the CDX index have gone into default. Yet, the pricing of these securities strongly suggests a substantial, and time-varying, fear of economic catastrophe.

The CDX and its tranches are an example of the structured finance products that proliferated in the period prior to the financial crisis of 2008–2009.[1] In the years following, both the academic literature and the popular press have deeply implicated structured finance in the series of events beginning with the near-default of Bear Stearns in March 2008 and culminating in collapse of Lehman brothers later that year.[2] Yet, despite the centrality of structured products to the crisis and its aftermath, there have been few attempts to quantitatively model these securities in a way that connects them to the underlying economy.

In this paper, we investigating whether an equilibrium model with rare economic disas-

[1] See Longstaff and Rajan (2008) for a description of structured finance products.

[2] See, for example, Reinhart and Rogoff (2009), Gorton and Metrick (2012), and Salmon (2009).

ters, in the spirit of Barro (2006) and Rietz (1988), can explain the time series of the cost to insure the CDX index and its tranches. Unlike previous quantitative models of structured finance, firm prices in our model are derived endogenously from assumptions on investor preferences and cash flows.[3] Firm prices embed rare disaster fears, as well as risks that are idiosyncratic. Importantly, our model can explain options as well as as equity prices (the model is also consistent with the average return and volatility on the aggregate market). We can therefore use it to back out a time series of rare disaster probabilities from option prices alone. When we use these probabilities to calculate model-implied values for CDX index and tranche prices, we find that it can explain the low spreads on senior tranches prior to the crisis, the high spreads during the crisis, and the timing of the increase in spreads. Our results imply that CDX spreads reflect an assessment of the risk in the economy that is consistent with other asset classes.

Our findings relate to a recent debate concerning the pricing of CDX tranches. Coval, Jurek, and Stafford (2009) examine the pre-crisis behavior of CDX tranches, pricing these tranches with a static options model that assumes that cash flows occur at a fixed maturity.[4] Setting firm parameters so that the CDX itself is correctly priced, they find that spreads for the senior tranches are too low in pre-crisis data. They conjecture that investors were willing to provide insurance on these products, despite receiving low spreads, because of a naive interpretation of credit risk ratings. Indeed, these products were highly rated because

[3]One strand of the credit derivative literature models default as an exogenous event which is the outcome of a Poisson process (Duffie and Singleton, 1997). Such models are known as "reduced form." A second strand builds on the assumption is that default occurs when firm value passes through a lower boundary (Black and Cox, 1976). Such models are known as "structural." Nonetheless, structural models typically take the price process and the risk-neutral measure as exogenous. In our paper, both are derived from investor preferences and beliefs.

[4]Specifically, if all cash flows occur at year five, then in principle CDX prices can be calculated using state prices derived from five-year options.

of their low default probabilities; these ratings did not take into account that defaults would occur during the worst economic states.

These conclusions are questioned by Collin-Dufresne, Goldstein, and Yang (2012). Collin-Dufresne et al. note that the pricing of the so-called equity tranche (the most junior tranche) is sensitive to the timing of defaults and the specification of idiosyncratic risk. Implicit in the pricing technique of Coval, Jurek, and Stafford (2009) is that default occurs at the five-year maturity. This need not be the case, and the difference could be important for the most junior tranches. Assuming that defaults occur at the five-year horizon makes the junior tranches look more attractive. The model spreads will thus be artificially low on the junior tranches, and, because the model is calibrated to match the index, the model-implied spreads on the senior tranches will be too high. Collin-Dufresne et al. also emphasize the need for fat-tails in the idiosyncratic risk of firms to capture the CDX spread at the three, as well as the five-year maturity. Introducing fat-tailed risk also raises the spread of junior tranches in the model and lowers the spread of senior tranches.

Once payoffs occur at a horizon other than five years, the method of extracting state prices from options data no longer cleanly applies. Accordingly, Collin-Dufresne, Goldstein, and Yang (2012) specify a dynamic model of the pricing kernel, which they calibrate to five-year options, and which they require to match three- and five-year CDX spreads. They find that this model comes closer to matching the spreads on equity tranches and on senior tranches prior to the crisis.

However, the results of Collin-Dufresne, Goldstein, and Yang (2012) point to the limitations of the methodology of both papers, in which asset prices are exogenous. Collin-Dufresne et al. show that pricing tranches during the crisis (October 2007 - September 2008) requires some probability of a catastrophic event that cannot be directly inferred from options data.[5]

[5]Without this catastrophic risk, matching the level of the CDX indices and option prices produces model-implied spreads on the senior tranches that are too *low*, while the junior tranche spreads are too high.

The probability of a catastrophic event cannot be determined by option prices using no-arbitrage arguments because there are not enough options with strikes in the relevant range. Thus CDX tranches are non-redundant securities. A model may explain option prices, but may fail to account for CDX senior tranches.

The limitations of the no-arbitrage framework lead us to an equilibrium model in which we derive firm valuations and the pricing kernel from assumptions on the endowment and investor preferences. Analytical solutions for firm prices and for options facilitate what would otherwise be an intractable numerical problem of computing CDX tranche prices. We require the resulting equilibrium model to explain the equity premium and equity volatility, and the low and smooth riskfree rate. Our model offers a joint, quantitative explanation for equity, options, and credit derivative pricing. Despite the constraints of the equilibrium approach, our model can match pre-crisis and crisis levels of the CDX and its tranches, ranging from equity to super-senior.

Our equilibrium model implies a link from rare disaster probabilities, to equity volatility, and from there to option prices. We can then use the time series of option prices to infer investor beliefs about rare event probabilities. Thus, besides matching the average levels, our results show that the same probabilities of rare events that price CDX senior tranches both before and during the crisis are fully consistent with prices on S&P index options. Moreover, the disaster probabilities implied by these prices are reasonable. Prior to the crisis, the disaster probability was close to zero. In September 2008 it rose precipitously, but the resulting level needed to explain CDX and CDX tranche prices is only 4% per annum. In our model, an shock to the probability of disaster endogenously lowers asset values, raises asset value volatility, and, of course, implies that future tail events are more likely. This combination of factors allows the model to match the level and time-variation in CDX/CDX

Matching the level of spreads during the crisis requires a lot of risk, and if this risk is idiosyncratic, it will lead to counterfactually high equity tranche spreads and counterfactually low senior tranche spreads.

tranche spreads.

Our results show that it is possible to account for the prices of senior tranches on the CDX within a frictionless model with reasonable parameter values. Moreover, the time series of these prices is consistent with the time series of options prices. We thus show that at least some of pricing behavior that was attributed to market failures during the crisis can be explained using the benchmark framework of representative agent asset pricing. Our findings also support the view that beliefs about rare disasters are an important determinant of stock market behavior.

The rest of this paper proceeds as follows. Section 2 describes the model and Section 3 describes the data. Section 4 demonstrates the virtual impossibility of describing these data using a lognormal model. This section argues that only a model with rare disasters would be capable of explaining CDX tranche prices. Section 5 describes the evaluation of our model using data on the aggregate market, on options and on CDX and CDX tranche data. Section 6 concludes.

2 Model

2.1 Model primitives and the state-price density

We assume an endowment economy with complete markets and an infinitely-lived representative agent. Aggregate consumption (the endowment) solves the following stochastic differential equation:

$$\frac{dC_t}{C_{t^-}} = \mu_c dt + \sigma_c dB_{c,t} + (e^{Z_{c,t}} - 1)dN_{c,t}, \tag{1}$$

where $B_{c,t}$ is a standard Brownian motion and $N_{c,t}$ is a Poisson process. The intensity of $N_{c,t}$ is given by λ_t and assumed to be governed by the following system of equations:

$$d\lambda_t = \kappa_\lambda(\xi_t - \lambda_t)dt + \sigma_\lambda\sqrt{\lambda_t}dB_{\lambda,t} \tag{2a}$$

$$d\xi_t = \kappa_\xi(\bar{\xi} - \xi_t)dt + \sigma_\xi\sqrt{\xi_t}dB_{\xi,t}, \tag{2b}$$

where $B_{\lambda,t}$ and $B_{\xi,t}$ are Brownian motions (independent of each other and of $B_{c,t}$).

The process defined by (1) has both normal-times risk, as represented by the Brownian component $\sigma_c dB_{c,t}$, and a risk of rare disasters represented by the Poisson term $(e^{Z_{c,t}} - 1)dN_{c,t}$. That is, at time t the economy will undergo a disaster with probability λ_t.[6] Given a disaster, the change in consumption (as a fraction the total) is $e^{Z_{c,t}} - 1$, where $Z_{c,t} < 0$ is a random variable. By writing the change in consumption as an exponential we ensure that consumption itself remains positive. For simplicity, we assume the distribution of $Z_{c,t}$ is time-invariant.

The system of equations (2) implies that the probability of a disaster λ_t is time-varying, and that it mean-reverts to a value ξ_t that itself changes over time. This type of multifrequency process has often been used for modeling asset price volatility and for option pricing in reduced-form models (see Duffie, Pan, and Singleton (2000) for discussion and references). Two-factor multifrequency processes are also used in the CDX literature (Collin-Dufresne, Goldstein, and Yang, 2012). Because return volatility will inherit the multifrequency variation of λ_t, (2a) is a natural choice for the disaster probability process. The process (2) can capture long memory in a time series, namely autocorrelations that decay at a slower-than-geometric rate. For example, the 2008 financial crisis was characterized both by a spike in λ_t that decayed quickly, and higher disaster probabilities in subsequent years. The equation system (2) captures this feature of the data, while a univariate autoregressive process would

[6] This description of Poisson shocks, which we adopt throughout the text, is approximate. In any finite interval there could theoretically be more than one shock since λ_t is an intensity, not a probability.

not. Setting σ_ξ to zero and assuming that ξ_t is at its (then deterministic) steady state of $\bar{\xi}$ results in the one-factor model of Wachter (2013) and Seo and Wachter (2016).

We assume a recursive generalization of time-additive power utility that allows for preferences over the timing of the resolution of uncertainty. Our formulation comes from Duffie and Epstein (1992), and we consider a special case in which the EIS is equal to one. That is, we define continuation utility V_t for the representative agent using the following recursion:

$$V_t = E_t \int_t^\infty f(C_s, V_s)\, ds,$$

where

$$f(C_t, V_t) = \beta(1-\gamma)V_t \left(\log C_t - \frac{1}{1-\gamma} \log((1-\gamma)V_t) \right).$$

The parameter β is the rate of time preference and γ is relative risk aversion. This utility function is equivalent to the continuous-time limit of the utility function defined by Epstein and Zin (1989) and Weil (1990). Assuming an EIS of one allows for closed-form solutions for equity prices up to ordinary differential equations, and facilitates the computation of options and CDX/CDX tranche prices.

In Appendix A, we show that the pricing kernel is characterized by the process

$$\frac{d\pi_t}{\pi_{t^-}} = \left(-r_t - \lambda_t E\left[e^{-\gamma Z_{c,t}} - 1\right]\right) dt \\ - \gamma\sigma_c dB_{c,t} + b_\lambda \sigma_\lambda \sqrt{\lambda_t} dB_{\lambda,t} + b_\xi \sigma_\xi \sqrt{\xi_t} dB_{\xi,t} + (e^{-\gamma Z_{c,t}} - 1) dN_{c,t}. \quad (3)$$

with

$$b_\lambda = \frac{\kappa_\lambda + \beta}{\sigma_\lambda^2} - \sqrt{\left(\frac{\kappa_\lambda + \beta}{\sigma_\lambda^2}\right)^2 - 2\frac{E\left[e^{(1-\gamma)Z_t} - 1\right]}{\sigma_\lambda^2}} \quad (4)$$

$$b_\xi = \frac{\kappa_\xi + \beta}{\sigma_\xi^2} - \sqrt{\left(\frac{\kappa_\xi + \beta}{\sigma_\xi^2}\right)^2 - 2\frac{b_\lambda \kappa_\lambda}{\sigma_\xi^2}}, \quad (5)$$

In the special case of time-additive utility, $\gamma = 1$ and $b_\lambda = b_\xi = 0$. The only risk that matters for computing expected returns is consumption risk, given by the terms $-\gamma\sigma_c dB_{c,t}$

and $(e^{-\gamma Z_{c,t}}-1)dN_{c,t}$, the latter of which captures the effect of rare disasters. For $\gamma > 1$ (this implies a preference for early resolution of uncertainty), b_λ and b_ξ are positive. Assets that increase in value when λ_t and ξ_t rise provide a hedge against disaster risk. All else equal, these assets will have lower expected returns, and higher prices, than otherwise.

The riskfree rate is given by

$$r_t = \beta + \mu_c - \gamma\sigma_c^2 + \lambda_t E\left[e^{(1-\gamma)Z_{c,t}} - e^{-\gamma Z_{c,t}}\right]. \tag{6}$$

Equation (6) implies that the riskfree rate is decreasing in the probability of an economic disaster. The greater is this probability, the more investors want to save for the future, and the lower the riskfree rate must be in equilibrium.

2.2 The aggregate market and index options

We assume that the aggregate market has payoff $D_t = C_t^\phi$ (Abel, 1990; Campbell, 2003). Empirically, dividends are more variable than consumption and more sensitive to economic disasters (Longstaff and Piazzesi, 2004). We capture this fact by setting $\phi > 1$. The process for dividends then follows from Ito's Lemma:

$$\frac{dD_t}{D_{t^-}} = \mu_d dt + \phi\sigma_c dB_{c,t} + (e^{\phi Z_{c,t}} - 1)dN_{c,t},$$

where $\mu_d = \phi\mu_c + \frac{1}{2}\phi(1-\phi)\sigma_c^2$.

In equilibrium, the price of the dividend claim is determined by the cash flows and the pricing kernel:

$$F(D_t, \lambda_t, \xi_t) = E_t\left[\int_t^\infty \frac{\pi_s}{\pi_t} D_s\, ds\right].. \tag{7}$$

Let $G(\lambda_t, \xi_t)$ be the ratio of prices to dividends. Then

$$\begin{aligned} G(\lambda_t, \xi_t) &= E_t\left[\int_t^\infty \frac{\pi_s}{\pi_t}\frac{D_s}{D_t}\, ds\right] \\ &= \int_0^\infty \exp\left(a_\phi(\tau) + b_{\phi\lambda}(\tau)\lambda_t + b_{\phi\xi}(\tau)\xi_t\right)\, d\tau, \end{aligned} \tag{8}$$

where $a_\phi(\tau), b_{\phi\lambda}(\tau)$ and $b_{\phi\xi}(\tau)$ satisfy ordinary differential equations given in Appendix B. Under the reasonable assumption of $\phi > 1$, $b_{\phi\lambda}(\tau) < 0$ (Wachter, 2013). Further, using the reasoning of Tsai and Wachter (2015), it can be shown that $b_{\phi\xi}(\tau) < 0$. The value of the aggregate market is thus decreasing in the disaster probability λ_t and its time-varying mean ξ_t.

Figure 1 shows the functions $b_{\phi\xi}(\tau)$ and $b_{\phi\lambda}(\tau)$, given the parameter values discussed in Section 5.1.1. Both functions are negative and decreasing as a function of τ. The function $b_{\phi\lambda}(\tau)$ converges after about 20 years reflecting the relative lack of persistence in λ_t. In contrast, the function $b_{\phi\xi}(\tau)$ takes nearly 70 years to converge. For dividend claims with maturities of 10 years or less, $b_{\phi\lambda}(\tau)$ is greater in magnitude than $b_{\phi\xi}(\tau)$, namely λ_t-risk is more important. However, in the limit as the horizon approaches infinity, the effect of ξ_t is nearly three times as large as the effect of λ_t.[7]

Applying Ito's Lemma to $F_t = D_t G(\lambda_t, \xi_t)$ gives the equilibrium law of motion for the aggregate market:

$$\frac{dF_t}{F_{t^-}} = \mu_{F,t}\,dt + \phi\sigma_c dB_{c,t} + \frac{\partial G}{\partial \lambda}\frac{1}{G}\sigma_\lambda\sqrt{\lambda_t}\,dB_{\lambda,t} + \frac{\partial G}{\partial \xi}\frac{1}{G}\sigma_\xi\sqrt{\xi_t}\,dB_{\xi,t} + \left(e^{\phi Z_{c,t}} - 1\right)\,dN_{c,t}, \quad (9)$$

The terms $\phi\sigma dB_{c,t}$ and $\left(e^{\phi Z_{c,t}} - 1\right)\,dN_{c,t}$ represent normal-time and disaster-time variation in dividends, respectively.[8] At our parameter values, the latter is a much more important source of risk than the former. Variation in λ_t and ξ_t produce variation in the price-dividend ratio G, and thus in stock prices. This is reflected in the terms $\frac{\partial G}{\partial \lambda}\frac{1}{G}\sigma_\lambda\sqrt{\lambda_t}\,dB_{\lambda,t}$ and $\frac{\partial G}{\partial \xi}\frac{1}{G}\sigma_\xi\sqrt{\xi_t}\,dB_{\xi,t}$. These can lead to highly volatile stock prices, even during normal times. Combining equa-

[7]See Lettau and Wachter (2007) and Borovička, Hansen, Hendricks, and Scheinkman (2011) for discussion of fixed-maturity dividend claims and the effect of exposures to risks of varying frequencies.

[8]The drift rate $\mu_{F,t}$ is determined in equilibrium from the instantaneous expected equity return r_t^e:

$$\mu_{F,t} = r_t^e - \lambda_t E\left[e^{\phi Z_{c,t}} - 1\right] + \frac{D_t}{F_t},$$

where r_t^e is determined by (10).

tions for the pricing kernel and for the aggregate market leads to the equation for the equity premium:

$$r_t^e - r_t = \gamma\phi\sigma^2 - \lambda_t E_t\left[\left(e^{-\gamma Z_{c,t}} - 1\right)\left(e^{-\phi Z_{c,t}} - 1\right)\right] - \lambda_t \frac{1}{G}\frac{\partial G}{\partial \lambda} b_\lambda \sigma_\lambda - \xi_t \frac{1}{G}\frac{\partial G}{\partial \xi} b_\xi \sigma_\xi \tag{10}$$

(Tsai and Wachter, 2015). The first term is from the CCAPM, and is negligible in our calibration. The next term is the risk premium due to disasters itself, and is positive and large. It represents the comovement of marginal utility and firm value during disaster times. The last two terms arise from time-variation in the risk of a disaster, both due to changes in λ_t and changes in ξ_t. Because disasters increase marginal utility (b_λ and b_ξ are positive), and decrease prices, these terms are positive. A calibration, like ours, that matches aggregate stock market volatility, also implies that they have a significant impact on the equity premium.

Our analysis below will rely on the prices of put options written on the aggregate market. A European put option gives the holder the right to sell the underlying security at some expiration date T for an exercise price K. Because the payoff on the option is $(K - F_T)^+$, no-arbitrage implies that

$$P(F_t, \lambda_t, \xi_t, T - t; K) = E_t\left[\frac{\pi_T}{\pi_t}\left(K - F_T\right)^+\right].$$

Lelt $K^n = K/F_t$ denote the normalized strike price ("moneyness") and $P_t^n = P_t/F_t$ the normalized put price. Like the price-dividend ratio, the normalized put price is a function of λ_t and ξ_t alone:

$$P^n(\lambda_t, \xi_t, T - t; K^n) = E_t\left[\frac{\pi_T}{\pi_t}\left(K^n - \frac{F_T}{F_t}\right)^+\right]. \tag{11}$$

We use (11) to calculate normalized put prices, and then find implied volatilities as defined by Black and Scholes (1973) (see Seo and Wachter (2016) for further details). As we show in Appendix G, the transform analysis of Duffie, Pan, and Singleton (2000) allows us to compute (11) analytically, avoiding the need for extensive simulations.

2.3 Individual firm dynamics

As explained in Section 2.4 below, CDX and CDX tranche pricing requires a model for individual firms. Let $D_{i,t}$ be the payout amount of firm i, for $i = 1, \cdots, N_f$, where N_f is the number of firms in the CDX (the number has been 125). While we use the notation $D_{i,t}$, we intend this to mean the payout not only to the equity holders but to the bondholders as well. The firm payout is subject to three types of risk:

$$\frac{dD_{i,t}}{D_{i,t^-}} = \mu_i dt + \underbrace{\phi_i \sigma_c dB_{c,t} + (e^{\phi_i Z_{c,t}} - 1) dN_{c,t}}_{\text{aggregate risk}} + \underbrace{I_{i,t}(e^{Z_{S_i}} - 1) dN_{S_i,t}}_{\text{sector risk}} + \underbrace{(e^{Z_i} - 1) dN_{i,t}}_{\text{idiosyncratic risk}}. \quad (12)$$

where μ_i is defined similarly to μ_d, namely $\mu_i = \phi_i \mu_c + \frac{1}{2}\phi_i(1 - \phi_i)\sigma_c^2$. The systematic risk is standard: D_{it} has a multiplicative component that behaves like $C_t^{\phi_i}$, analogously to dividends. Firms are exposed to both normal-times aggregate risk and aggregate consumption disasters. Because financial leverage is not reflected in C_t, the value of ϕ_i will be substantially below that of ϕ for aggregate (equity) dividends above. However, we will still allow firms to have greater exposure to aggregate disasters than consumption, namely $\phi_i > 1$ (labor income could account for the wedge between unlevered cash flows and consumption).

Firms are also exposed to idiosyncratic negative events that occur with constant probability λ_i. For simplicity, we assume all idiosyncratic risk is Poisson.[9] When a firm is hit by its idiosyncratic shock (which we model as an increment to the counting process $N_{i,t}$), the firm's payout falls by $D_{i,t^-} \times (1 - e^{Z_i})$. For parsimony, we assume μ_i, σ_i, ϕ_i, Z_i, and λ_i are the same for all i, and that Z_i is a single value (rather than a distribution). The shocks themselves, $dN_{i,t}$, will of course be independent of one another and independent of the aggregate shock $dN_{c,t}$. Longstaff and Rajan (2008) estimate that a portion of the CDX spread is attributable to risk that affects a nontrivial subset of firms. Following Longstaff

[9] Campbell and Taksler (2003) show that idiosyncratic risk and the probability of firm default are strongly linked in the data.

and Rajan and Duffie and Garleanu (2001), we refer to this as sector risk, and allow for it in (12). Let $\mathcal{S}$ denote a finite set of sectors.[10] Each firm is in exactly one sector; we let $S_i \in \mathcal{S}$ denote the sector for firm i and $dN_{S_i,t}$ the sector shock. When a sector shock arrives, the firm is hit with probability p_i, namely the sector term in (12) is multiplied by $I_{i,t}$ which takes a value 1 with probability p_i and 0 otherwise. If a firm happens to be affected by this sector shock, the firm's payout drops by $D_{i,t^-} \times (1 - e^{Z_{S_i}})$.[11] Again, for parsimony, p_i and Z_{S_i} are the same across firms. The shocks $I_{i,t}$ are independent across firms and $dN_{S_i,t}$ are independent across sectors.

Intuitively, sector risk should be correlated with aggregate consumption risk. To capture this correlation, we allow the intensity of $N_{S_i,t}$, $\lambda_{S_i,t}$, to depend on the state variables λ_t and ξ_t. In the Appendix, we solve for firm values under the specification $\lambda_{S_i,t} = w_0 + w_\lambda \lambda_t + w_\xi \xi_t$. For parsimony, we will calibrate the simpler model $\lambda_{S_i,t} = w_\xi \xi_t$.

Given this payout definition, we solve for the total value of firm i (the equity plus the debt), which we denote $A_i(D_{i,t}, \lambda_t, \xi_t)$:

$$A_i(D_{i,t}, \lambda_t, \xi_t) = E_t \left[\int_t^\infty \frac{\pi_s}{\pi_t} D_{i,s} ds \right]. \qquad (13)$$

Define the price-to-payout ratio as

$$G_i(\lambda_t, \xi_t) \equiv \frac{A_i(D_{i,t}, \lambda_t, \xi_t)}{D_{i,t}}.$$

Similar to the price-dividend ratio, the price-to-payout ratio for an individual firm can be

[10]For concreteness, we can think of the sector classification as corresponding to that given by our data provider Markit. There are five sectors: Consumer, Energy, Financials, Industrial, Telecom, Media, and Technology, so $\mathcal{S} = \{\text{C, E, F, I, T}\}$. We will use this classification to discipline our calibration. However, neither the equations nor our empirical results require this interpretation.

[11]This structure is isomorphic to one in which the set of firms is partitioned into a greater number of sectors and firms are hit by sector shocks with probability one.

expressed as an integral of an exponential-linear function of the state variables:

$$G_i(\lambda_t, \xi_t) = \int_0^\infty \exp\left(a_i(\tau) + b_{i\lambda}(\tau)\lambda_t + b_{i\xi}(\tau)\xi_t\right) d\tau, \tag{14}$$

where $a_i(\tau)$, $b_{i\lambda}(\tau)$, and $b_{i\xi}(\tau)$ solve the system of ordinary differential equations derived in Appendix C. Like the aggregate market value, firm values are decreasing in the disaster probability λ_t and its time-varying mean ξ_t.

The dynamics of firm values $A_{it} = A_i(D_{i,t}, \lambda_t, \xi_t)$ follow from Ito's Lemma:

$$\frac{dA_{i,t}}{A_{i,t^-}} = \mu_{A_i,t}\, dt + \phi_i \sigma_c dB_{c,t} + \frac{\partial G_i}{\partial \lambda}\frac{1}{G_i}\sigma_\lambda \sqrt{\lambda_t}\, dB_{\lambda,t} + \frac{\partial G_i}{\partial \xi}\frac{1}{G_i}\sigma_\xi \sqrt{\xi_t}\, dB_{\xi,t} + \left(e^{\phi_i Z_{c,t}} - 1\right) dN_{c,t} + I_{i,t}(e^{Z_{S_i}} - 1)dN_{S_i,t} + (e^{Z_i} - 1)dN_{i,t} \tag{15}$$

where $\mu_{A_i,t}$ is the asset drift rate, determined in equilibrium. Equation 15 has some similarity to processes that use reduced-form models for asset prices. For example, as in Collin-Dufresne, Goldstein, and Yang (2012), there is Brownian risk with stochastic volatility (following a multifrequency process), a risk of an adverse idiosyncratic event, and a risk of catastrophic market-wide decline. Here, however, the process is an endogenous outcome of our assumptions on fundamentals and on the utility function. Specifically, volatility occurs because of changes in agents' rational forecasts of economic disasters.

2.4 CDX pricing

A credit default swap contract provides a means of trading on the risk of default of a single firm. Under this bilateral contract, the protection buyer commits to paying an insurance premium to the protection seller, who pays the protection buyer the loss amount in the case of default. Recent statistical models of single-name credit default swaps suggest that market participants price in the risk of rare idiosyncratic and market-wide events (Kelly, Manzo, and Palhares, 2016; Seo, 2014). Our focus in this paper is on the CDX North

American Investment Grade, an index whose value is determined by default events on a set of underlying firms, also known as reference entities. An investor who buys protection on the CDX is, in effect, buying protection on default events of all the underlying firms.

At each time t, we price a CDX contract initiated at t and maturing at T on N_f reference entities whose asset prices are governed by (15). Define default as the event that a firm's value falls below a threshold A_B (Black and Cox, 1976). The default time for firm i is therefore

$$\tau_{t,i} \quad = \quad \inf\left\{\tau > t : \frac{A_i(D_{i,\tau}, \lambda_\tau, \xi_\tau)}{A_i(D_{i,t}, \lambda_t, \xi_t)} \leq A_B\right\}. \tag{16}$$

To maintain stationarity, we define the threshold relative to the value of the firm at the initiation.[12] Let $R_{\tau_{t,i}}$ denote the recovery rate for a firm defaulting at $\tau_{t,i}$. This recovery rate is a random variable that depends only on the outcome of $dN_{c,\tau_{t,i}}$, namely whether default co-occurs with a disaster. Note that this specification implies that the time-t distribution of both $\tau_{t,i} - t$ and of $R_{\tau_{t,i}}$ is completely determined by λ_t and ξ_t.

We first discuss contracts on CDX as a whole, and then consider tranches in Section 2.5. Following the convention in our data, we assume that the protection buyer is paying to insure \$1 (namely, \$1 is the notional). If firm i defaults, the loss on the CDX increases by $\frac{1}{N_f}(1 - R_{\tau_{t,i}})$. Let $L_{t,s}$ denote the cumulative loss at time s. Then $L_{t,s}$ is given by

$$L_{t,s} = \frac{1}{N_f}\sum_{i=1}^{N_f} 1_{\{t<\tau_{t,i}\leq s\}}(1 - R_{\tau_{t,i}}). \tag{17}$$

Increases in $L_{t,s}$ trigger payments from the protection seller (the party providing insurance)

[12] If we interpret $A_B A_{i,t}$ as the amount of debt that the firm has at time t, then (16) implies that the firm is in default whenever the value of equity is below zero. This formulation is consistent with firms maintaining stationary leverage, and but that reversion to the stationary levels takes place on the order of years (Lemmon, Roberts, and Zender, 2008).

to the protection buyer. No-arbitrage implies that the value of these payments equals

$$\textbf{Prot}_{\text{CDX}}(\lambda_t, \xi_t; T-t) = E_t^Q \left[\int_t^T e^{-\int_t^s r_u du} dL_{t,s} \right],$$

where E^Q denotes the expectation taken under the risk-neutral measure Q and r_u is the riskfree rate, both of which are implied by the pricing kernel (3).[13] Equation 18 is sometimes referred to as the "protection leg" of the contract.

In a CDX contract, the protection buyer makes payments at quarterly intervals. If no firms default, these premium payments add up to a fixed value S over the course of a year. If default occurs, the premium payments fall to reflect the fact that a lower amount is now insured under the contract. Let $n_{t,s}$ denote the fraction of firms that have defaulted $s-t$ years into the contract:

$$n_{t,s} = \frac{1}{N_f} \sum_{i=1}^{N_f} 1_{\{t < \tau_{t,i} \le s\}}. \tag{18}$$

Like $L_{t,s}$, $n_{t,s}$ is a random variable whose value is realized at s and whose time-t distribution depends only on λ_t, ξ_t, and $s-t$. Let $\vartheta = 1/4$, the interval between premium payment dates. For a given spread S, the premium leg is equal to

$$\textbf{Prem}_{\text{CDX}}(\lambda_t, \xi_t; T-t, S) = S E_t^Q \Bigg[\frac{1}{4} \sum_{m=1}^{4T} e^{-\int_t^{t+\vartheta m} r_u du} (1 - n_{t,t+\vartheta m}) + \int_{t+\vartheta(m-1)}^{t+\vartheta m} e^{-\int_t^s r_u du} (s - t - \vartheta(m-1)) dn_{t,s} \Bigg]. \tag{19}$$

The first term in (19) is the value of the scheduled premium payments. Note that $n_{t,t+\vartheta m}$ is the fraction of the pool that has defaulted as of the mth payment, and so $1 - n_{t,t+\vartheta m} \le 1$ is the notional at that point in time. The second term represents the accrued premium.

[13] Following longstanding practice in the literature on credit derivatives, we express prices as discounted cash flows under the risk neutral measure rather than as cash flows multiplied by π_t under the physical measure. The mapping between the two is well-known (Duffie, 2001).

Consider the event of firm default at some time s between payments $m-1$ and m. Until time s, the protection buyer is insured against default of this firm. However, this is not reflected in the scheduled payments; the value of the mth payment is the same as it would be if the default occurred at time $t+\vartheta(m-1) < s$. Thus upon default the protection buyer pays the fraction of $S/(4N_f)$ that accrues between the $(m-1)$th payment and the default time s, known as the accrued premium.

The CDX spread $S_{\text{CDX}}(\lambda_t, \xi_t; T-t)$ is the value of S that equates the premium leg (19) with the protection leg (18). We describe the computation of this spread in Appendix D.

2.5 CDX tranche pricing

Generally, tranches are claims that partition a security and have different levels of subordination in case of default. In the case of the CDX (which is a synthetic collateralized debt obligation and thus does not represent underlying physical assets), tranche losses are defined in terms of the total loss $L_{t,s}$.

Tranches are defined by two numbers: the attachment point, which gives the level of CDX loss at which the tranche is penetrated, and the detachment point, after which further losses detach from the tranche. For example, consider the 10-15% tranche. This tranche loses value if the CDX accumulates more than a 10% loss. Losses of between 10% and 15% attach to this tranche. If the losses reach 15%, the notional amount of the tranche is exhausted. Further losses attach to the next tranche. During our sample period, six tranches commonly traded, with attachment-detachment pairs 0-3%, 3-7%, 7-10%, 10-15%, 15-30%, and 30-100%. The most junior tranche (0-3%) is referred to as equity, the second as mezzanine, and the remaining four as senior. The last, and most senior tranche (30-100%), is often called "super senior."

Let K_{j-1} be the attachment point and K_j the detachment point of the j-th tranche for

$j = 1, \ldots, J$, where $K_0 = 0$ and $K_J = 1$. Given a CDX loss $L_{t,s}$, the tranche loss is given by

$$T^L_{j,t,s} = T^L_j(L_{t,s}) = \frac{\min\{L_{t,s}, K_j\} - \min\{L_{t,s}, K_{j-1}\}}{K_j - K_{j-1}}. \tag{20}$$

If the CDX loss is below both K_j and K_{j-1}, it has not attached to the tranche, and the loss is 0%. If the loss is greater than both K_j and K_{j-1}, it detaches from the tranche and the tranche loss is 100%. If the loss is between K_{j-1} and K_j, then the loss equals $0 < \frac{L_{t,s} - K_{j-1}}{K_j - K_{j-1}} < 1$. The definition (20) implies that the notional amount on each tranche is equal to \$1, which is the convention in our data. Note that the weighted sum of tranche losses equals the total loss:

$$\sum_{j=1}^{J} (K_j - K_{j-1}) T^L_{j,t,s} = L_{t,s}. \tag{21}$$

Given the specification for the tranche loss, the protection seller for tranche j pays

$$\mathbf{Prot}_{\text{Tran},j}(\lambda_t, \xi_t, T - t) = E^Q_t \left[\int_t^T e^{-\int_t^s r_u du} dT^L_{j,t,s} \right]. \tag{22}$$

The protection buyer for tranche j makes quarterly premium payments. As in the case of the CDX, the amount he or she pays depends on the tranche notional. However, the adjustments in tranche notional for default is more complicated than for the CDX. The adjustment in notional depends not only on the tranche loss, but also on something called tranche recovery. If a firm defaults, the notional on the CDX falls by $\frac{1}{N_f}$. However, the notional on the most junior tranche falls by the smaller amount of $\frac{1}{N_f}(1 - R_{\tau_{t,i}})$.[14] To keep the notional amount on the CDX tranches consistent with that of the CDX, the total change in notional on the tranches following default must also add up to $\frac{1}{N_f}$. This extra reduction in notional is called tranche recovery.

It is customary to apply the tranche recovery to the most senior tranche. Note that $n_{t,s} - L_{t,s}$ is the amount recovered to date from defaults. Then tranche recovery for the

[14] To be precise, this is the change in notional multiplied by the width of the tranche, $K_j - K_{j-1}$.

super-senior tranche defined to be

$$T^R_{J,s} = \frac{n_{t,s} - L_{t,s}}{K_J - K_{J-1}}$$

(recall that $K_J = 1$). In the very rare event that this recovery exhausts the notional on the senior tranche, the remaining recovery amount detaches from this tranche and attaches to the next most senior tranche. A general definition for tranche recovery is thus

$$T^R_{j,t,s} = T^R_j(n_{t,s} - L_{t,s}) = \frac{\min\{n_{t,s} - L_{t,s}, 1 - K_{j-1}\} - \min\{n_{t,s} - L_{t,s}, 1 - K_j\}}{K_j - K_{j-1}}. \tag{23}$$

Note that if $n_{t,s} - L_{t,s} < 1 - K_j$, then no recovery applies to tranche j (it can all be applied to the more senior tranches). It follows that $n_{t,s} - L_{t,s} < 1 - K_{j-1}$ as well, and (23) is equal to zero. If on the other hand, $n_{t,s} - L_{t,s}$ is greater than $1 - K_j$ and less than $1 - K_{j-1}$, then the numerator in (23) equals $n_{t,s} - L_{t,s} - (1 - K_j)$, which is the amount of the recovery not reflected in the tranche loss for more senior tranches. If $n_{t,s} - L_{t,s}$ is above both $1 - K_j$ and $1 - K_{j-1}$ then (23) is equal to 1. As is the case for tranche losses (24), the weighted sum of tranche recovery is equal to total recovery:

$$\sum_{j=1}^{J} (K_j - K_{j-1}) T^R_{j,t,s} = n_{t,s} - L_{t,s}. \tag{24}$$

Combining (21) and (24), we see that the weighted change in notional from a default is $n_{t,s}$, which is the change in notional for a contract on the CDX itself.

Given these definitions of the tranche loss and recovery, we can define the premium payments for a given tranche. Let S be the spread, and U the upfront payment (to be discussed further below). Then the premium leg for tranche j is given by

$$\mathbf{Prem}_{\mathrm{Tran},j}(\lambda_t, \xi_t; T - t, U, S) = \\ U + SE^Q \left[\frac{1}{4} \sum_{m=1}^{4T} e^{-\int_t^{t+\vartheta m} r_u du} \int_{t+\vartheta(m-1)}^{t+\vartheta m} \left(1 - T^L_{j,t,s} - T^R_{j,t,s}\right) ds \right] \tag{25}$$

(the definitions of tranche loss and recovery imply that their sum cannot exceed 1). Except for the upfront payment U, the tranche premium leg is nearly equivalent to that of the CDX, (19), as can be shown by integration by parts. The difference is in the timing of the accrued premium payment. For the CDX, this payment is made upon occurrence of default. For the tranche, this payment is made at the next scheduled premium payment date.[15] The difference in the contract terms may reflect the fact that the event that a loss or recovery attaches to a tranche is more difficult to determine than a default event, because it requires the computation of $L_{t,s}$ rather than $n_{t,s}$. The computation of $L_{t,s}$ requires firm i's recovery, $R_{\tau_{t,i}}$, which may depend on the outcome of ISDA Credit Event Auction proceedings, and thus will not typically be known in real time.

For all but the equity tranche, the upfront payment U is set to zero in our data, and the spread S is determined in the same way as the CDX as a whole. However, the equity tranche trades assuming a set spread of 500 basis points, with U determined so as to equate the premium and the protection legs. Why this difference? As Collin-Dufresne, Goldstein, and Yang (2012) point out, the spread payments for the equity tranche are particularly sensitive to losses, both because equity is hit first when default occurs, and because the CDX loss is likely to be large relative to the detachment point of 3%. In practice, it is difficult to know the precise moment of default, which affects the amount of spread to be paid. Setting an upfront payment reduces the spread, and therefore reduces the sensitivity of the cash flows to the precise timing of a default event.

Following the conventions in the data, therefore, define $U_{\text{Tran},1}(\lambda_t, \xi_t; T-t)$ to be the

[15] Recognizing that $n_{t,s}$ is the CDX equivalent of $T^L_{j,t,s} + T^R_{j,t,s}$, the CDX equivalent of the integral in (25) is

$$\int_{t+\vartheta(m-1)}^{t+\vartheta m} (1-n_{t,s})\, ds = \frac{1}{4}(1-n_{t,t+\vartheta m}) + \int_{t+\vartheta(m-1)}^{t+\vartheta m} (s-t-\vartheta(m-1))\, dn_{t,s},$$

where the right hand side is derived using integration by parts. Note that this right hand side is equivalent to the change in notional in (19), except for a change in the discounting of the accrued premium payment.

value of U that equates (25) and (22), for $j = 1$, and $S = 0.05$. For $j = 2, \ldots, J$, define $S_{\text{Tran},j}(\lambda_t, \xi_t; T - t)$ to be the value of S that, for $U = 0$ equates (25) with (22). Appendix D describes CDX tranche pricing.

3 Data

Our analyses require the use of pricing data from options and CDX markets. The options data, provided by OptionMetrics, consist of daily implied volatilities on S&P 500 European put options from January 1996 to December 2012. To construct a monthly time series, we use data from the Wednesday of every option expiration week. We apply standard filters to extract contracts with meaningful trade volumes and prices. To obtain an implied volatility curve for each date, we fit a polynomial in strike price and maturity (see Seo and Wachter (2016) for more details on options data construction).

Our CDX data come from Markit and consist of daily spreads and upfront amounts for the October 2005 to September 2008 period on the 5-year CDX North American Investment Grade index and its tranches excluding the super-senior. The CDX North American Investment Grade index is the most actively traded CDX product. We will refer to it in what follows as "the CDX." This index represents 125 equally-weighted large North American firms that are investment-grade at the time the series is initiated.

To maintain an (approximately) fixed-maturity contract, a new series for the CDX is introduced every March and September and the previous series becomes off-the-run. Our sample corresponds to CDX series 5 through 10.[16] We use data from the series that is

[16] The liquidity of CDX tranches significantly shrank after CDX10, the last series introduced before the Lehman crisis. From series 11 on, these products were traded too infrequently for prices to be meaningful. Moreira and Savov (2016) presents a model with time-varying disaster risk that accounts for qualitative features of structured finance around the crisis, including the lack of trading in these securities following the

most recently issued, and hence most actively traded, in our analyses. For comparability with prior studies (Collin-Dufresne, Goldstein, and Yang, 2012; Coval, Jurek, and Stafford, 2009), we report average spreads for two subperiods, with September 2007 being the end of the first subperiod. In our sample, the CDX and all tranches except for the equity tranche are quoted in terms of spreads. The equity tranche is quoted in terms of upfront payment with a fixed spread of 500 basis points.[17]

4 Why disaster risk?

Before quantitatively evaluating our model, we motivate our approach by examining what would happen under a lognormal distribution for asset values.[18] We will concern ourselves in this section with the two most senior tranches, implying that we are interested in losses of 15% or more on the CDX. Given the effects of diversification, it is very unlikely for this level of losses to occur due to anything other than a severe market-wide shock. It therefore suffices to consider a single firm whose value follows a geometric Brownian motion. For simplicity, we assume zero recovery and ignore discounting.[19] Under these assumptions, the value of the protection leg (22) equals the default probability of this single firm calculated under the risk-neutral measure.

The assumption of a geometric Brownian motion implies that the change in firm value

Lehman default.

[17]Trading conventions changed with the introduction of the Standard North American Contract (SNAC) in April 2009. Under SNAC, CDX products trade with upfront amounts and fixed coupons of either 100 or 500 bps.

[18]To make this argument as transparent as possible, we assume that shocks are homoskedastic. However, there is a strong reason to think that our results will generalize to the heteroskedastic case: Collin-Dufresne, Goldstein, and Yang (2012) show that a model with stochastic volatility but no Poisson shocks generates strikingly counterfactual predictions when forced to fit tranche data.

[19]Allowing for recovery will make it harder for the lognormal model to fit the data.

between time s and time t is given by

$$\log A_s - \log A_t = \mu_A(s-t) + \sigma_A(B_s - B_t),$$

where B_t is a standard Brownian motion under the risk-neutral measure. To give the lognormal model its best chance of success, we consider values for σ_A that are higher than what data would suggest, ranging from 14% to 20% (under lognormality, the physical and risk-neutral volatilities are identical). Note that these volatilities are for equity plus debt. Historical stock return volatility is 18%; assuming leverage of 32% (see below) asset return volatility would be 12%. Another useful benchmark is what our model, calibrated to option prices, says about firm volatility during the crisis period. Substituting in the average value of λ_t and ξ_t over the crisis subsample into (15), we find a volatility of 15%. However, using this value for σ_A implies a belief on the part of market participants that this high value is likely to persist for 3 to 5 years, which seems unlikely given the lack of persistence of volatility in the data.[20]

The probability of crossing the default boundary also depends on the drift. No-arbitrage implies that this drift is equal to $\mu_A = r - \delta - \frac{1}{2}\sigma_A^2$, where δ is the ratio of dividends plus interest to total firm value. We compute the probabilities of default for values ranging from -6% to -2%. While these numbers are low given historical data, higher values lead to lower probabilities, so choosing low numbers gives the lognormal model its best chance at success.

Table 4 shows the resulting probability of default over three and five-year periods, defined

[20]Note that accounting for this level of volatility given fundamentals is not an easy hurdle in and of itself. Models that have the potential to do so (besides the time-varying disaster risk model that is the focus of this paper) include Bansal and Yaron (2004) and Campbell and Cochrane (1999). These models are conditionally lognormal. Returns over multiple periods will not be lognormal due to stochastic volatility. However, as noted above, non-normalities due only to stochastic volatility are very unlikely to explain senior tranche spreads.

as firm value below the boundary A_B.[21] As the previous discussion indicates, we can think of these probabilities as approximating the value of the protection leg on the tranches. Because the annual spread equates the protection leg and the premium leg, the spread should (again, roughly), equal the probability divided by the number of years in the contract.

Table 4 shows that, for a drift of -6% and a volatility of 20%, the probability of default in three years or less is 0.0041%. That is, there is a less than one in ten thousand chance that default would occur in under three years. The average annual spread for 3-year contracts is 48 basis points for the third senior tranche and 23 basis points for the super-senior tranche. Even ignoring the fact that we should multiply these tranche spreads by three to obtain the total payment, they are two orders of magnitude too high, given the predictions of the lognormal model.

While the failure of the lognormal model is most dramatic for 3-year contracts, we can also see it in spreads on 5-year contracts. The highest probability shown in Table 4, corresponding to very high asset value volatility and low drift, is 0.43% over five years. While this is the same order of magnitude as the 5-year spread on the third senior tranche (69 basis points), keep in mind that the 5-year spread is annual, and should be multiplied by 5 to compare to the probability. Thus 5-year probabilities for the lognormal distribution are also clearly unrealistic.[22] These calculations suggest that only a model that admits large, sudden, and

[21] The distribution of the default time τ defined in (16) is given by

$$P(\tau < u) = 1 - \Phi\left(\frac{-\log A_B + \mu_A}{\sigma_A \sqrt{u}}\right) + e^{\frac{\mu_A \log A_B}{\sigma_A^2}} \Phi\left(\frac{-\log A_B - \mu_A}{\sigma_A \sqrt{u}}\right),$$

where $\Phi(\cdot)$ is the normal cumulative distribution function.

[22] Nonetheless, the short-term CDX senior tranche spreads are particularly powerful in ruling out a log-normal model. This result has a parallel in the literature on single-name corporate bonds. Zhou (2001) shows that fat-tailed idiosyncratic risk is necessary to match single-name short-term bond yields. Culp, Nozawa, and Veronesi (2014) show that securities that are equivalent, by no-arbitrage, to even shorter-term bonds also have high credit spreads. Their results are further evidence of the need for fat-tailed events that

pervasive declines in firm values can explain the level of CDX tranche spreads during the crisis. Such declines must be rare because such an episode has not been observed within the last 100 years of U.S. history.

The analysis in this section pertains to the risk-neutral process for prices, rather than a physical process for cash flows and for the pricing kernel. That is, while the argument establishes that the risk-neutral process for asset values must contain catastrophic jumps, it does not imply that rare disasters are a feature of the physical world, nor that they have an important affect on risk premia. Is this a limitation? We argue that it is not. If the risk-neutral process for prices has jumps, the physical process must too (because of the equivalence of the measures). Unless risk aversion is extraordinarily large, the jumps in the physical process must also be catastrophic. Once the physical process has large jumps in prices, such large jumps cannot but play an important role in determining risk premia, both for individual firms and for the market as a whole, as we can see from the jump component in the equation for risk premia, (10). Constant relative risk aversion implies that large declines are much more costly for the representative agent then small declines, and the declines due to catastrophic shocks implied by CDX tranche prices are quite large.

What this argument does not identify is the source of the decline in prices. The model in Section 2 assumes an instantaneous permanent decline in consumption and firm cash flows. A model with Poisson shocks to consumption and dividend drifts generates similar results (Tsai and Wachter, 2015), as might a model in which volatility jumps upward and is persistent.[23] What is necessary is that some aspect of the distribution of firm fundamentals can shift suddenly and unpredictably in an unfavorable way, and that these shifts are far

are economy-wide.

[23]When disasters result in a instantaneous decline in consumption and when dividends are given by $D_t = C_t^{\phi}$, $e^{-\phi Z_c} - 1$ is the change in value of the market portfolio. In general, however, Poisson shocks to the distribution of dividends will be reflected in large changes market values.

beyond what one would expect from a normal distribution.

5 Evaluating the model

A standard approach to comparing an endowment economy model with the data is to simulate population moments and compare them with data moments. In a model with rare disasters, this may not be the right approach if one is looking at a historical period that does not contain a disaster. An alternative approach is to simulate many samples from the stationary distribution implied by the model, and see if the data moments fall between the 5th and 95th percentile values simulated from the model. For this study, this approach is not ideal for two reasons. First, the short length of CDX/CDX tranche time series will likely mean that the error bars implied by the model will be very wide. Thus this test will have low power to reject the model. Second, unlike stock prices which are available in semi-closed-form, and options, which are available up to a (hard-to-compute, but nonetheless one-dimensional) integral, CDX prices must be simulated for every draw from the state variables. Thus the simulation approach is computationally infeasible.

For these reasons, we adopt a different approach. We first extract the time series of our two state variables from options data and then generate predictions for the CDX index and tranches based on this series of state variables. We are thus setting up a more stringent test than endowment economy models are usually subject to. Namely, we are asking that the model match not only moments, but the actual time series of variables of interest.

5.1 Calibration

In this section, we describe how we calibrate the model. Section 5.1.1 describes the choice of parameter values for the utility, consumption, and dividend processes, as well as the fit to moments of the aggregate market and the riskfree rate in the data. Section 5.1.2 describes the calibration of firm-level parameters.

5.1.1 Aggregate consumption and dividends

We keep risk aversion γ, the discount rate β, the leverage parameter ϕ, the normal-times consumption distribution, and the disaster distribution the same as in Wachter (2013) and Seo and Wachter (2016).[24] Note that κ_λ and σ_λ will not have the same interpretation in this model as in the earlier one.

Our first goal in calibrating the model is to generate reasonable predictions for the aggregate market and for the consumption distribution. One challenge in calibrating representative agent models is to match the high volatility of the price-dividend ratio. There is an upper limit to the amount of volatility that can be assumed in the state variable before a solution for utility fails to exist. The more persistent the processes, namely the lower the values of κ_λ and κ_ξ, the lower the respective volatilities must be so as to ensure that the discriminants in (4) and (5) stay nonnegative. We choose parameters so that the discrimant is equal to zero. Thus there are only a total of three free parameters to match the aggregate market, the riskfree rate, and consumption.

The resulting parameter choices are shown in Table 1. The mean reversion parameter κ_λ and volatility parameter σ_λ are relatively high, indicating a fast-moving component to the λ_t process, while the mean reversion parameter κ_ξ and σ_ξ are relatively low, indicating a slower-moving component. The parameter $\bar{\xi}$ (which represents both the average value of ξ_t and the average value of λ_t) is 2% per annum. This is a lower average disaster probability than in Wachter (2013), and, in this sense, our calibration is conservative. The extra persistence

[24] This disaster distribution is drawn from Barro and Ursúa (2008) who build on work of Barro (2006). We also assume a 40% probability of default on the government bill in the case of disaster as in Barro (2006) and following these two papers.

created by the ξ_t process implies that λ_t can deviate from its average for long periods of time. To clarify the implications of these parameter choices, we report population statistics on λ_t in Panel C of Table 1. The median disaster probability is only 0.37%, indicating a highly skewed distribution. The standard deviation is 3.9% and the monthly first-order autocorrelation is 0.986.

Implications for the riskfree rate and the market are shown in Table 2. We simulate 100,000 samples of length 60 years to capture features of the small-sample distribution (Appendix E describes to simulate from the model). We also simulate a long sample of 600,000 years to capture the population distribution. Statistics are reported for the full set of 100,000 samples, and the subset for which there are no disasters (38% of the sample paths). To be conservative, we take the view that this no-disaster distribution is the more appropriate point of comparison for postwar data. However, whether the 2008-2009 crisis constitutes a disaster in the U.S. depends on the data series one looks at (industrial production fell by more than consumption). As in earlier work (Wachter, 2013), time-varying disaster probability implies a high equity premium, low riskfree rate, and high equity volatility.[25] Only the autocorrelation of the price-dividend ratio falls (just) outside the 90% confidence bounds.[26]

[25] The average Treasury Bill rate is slightly too high, though this could be lowered by lowering β or by lowering the probability of government default.

[26] While it is possible to calibrate the model to match this autocorrelation, it comes at the cost of raising the autocorrelation of option prices beyond realistic levels. Moreover, so that utility converges, there is a tradeoff between persistence and volatility. One view is that the autocorrelation of the price-dividend ratio observed in the postwar period may be driven by very long-run fluctuations in the tendency to pay dividends (Fama and French, 2001), which are outside of the scope of the model. Note that these fluctuations could also be partially responsible for the observed volatility of the price-dividend ratio, which is also somewhat higher in the data than the model.

5.1.2 Firm-level payouts

To go from state variables to CDX/CDX tranche prices requires not only dynamics for the market, but dynamics for individual firms and assumptions about what constitutes a default. For parsimony, and following standard practice in this literature, we assume firms are ex ante identical. However, firms face distinct idiosyncratic shocks and potentially distinct sector-wide shocks.

To calibrate individual firm dynamics, we use results from Collin-Dufresne, Goldstein, and Yang (2012). The default boundary A_B is set to be 19.2%. This implies that if the asset value falls below 0.192 multiplied by what it was at the initiation of the contract, the firm is assumed to be in default. This value derives from the average leverage ratio from firm-level data (32%). Firms are then considered to be in default if their value is 60% of their debt outstanding.[27] We assume that the recovery rate is 40% in normal times and 20% in the event of rare disasters. We choose the idiosyncratic jump size $e^{Z_i} - 1$ to equal -80%, a value sufficiently large to make default almost certain. Both assumptions are consistent with those of Collin-Dufresne et al.

The CDX index consists of investment-grade firms that are relatively large and stable. Collin-Dufresne, Goldstein, and Yang (2012) estimate a different asset beta for each CDX series; the asset betas are between 0.5 and 0.6 for pre-crisis series and between 0.6 and 0.7 for crisis series. This reflects a slight increase in leverage for the firms included in the series. In our model, the asset beta will be mostly determined by the ratio ϕ_i/ϕ; however, the connection between sector-wide and aggregate risk adds an additional degree of covariance.

[27] The fact that our model can generate realistic CDX spreads given conservative assumptions on leverage ratios and default suggests that our model can resolve the credit spread puzzle, namely, it will be able to explain high credit spreads given low historical default rates. Papers that examine credit spreads from a disaster-risk perspective include Christoffersen, Du, and Elkamhi (2016), Gabaix (2012), and Gourio (2013). These papers do not look at the CDX.

We therefore choose $\phi_i/\phi = 0.5$ pre-crisis and $\phi_i/\phi = 0.6$ during the crisis. Given our assumption of $\phi = 2.6$, this corresponds to a pre-crisis value of ϕ_i of 1.3 and a crisis value of 1.6.[28]

We use the results of Longstaff and Rajan (2008) to calibrate the parameters for the sector-wide shocks. For simplicity, in the calibration we assume that the loss in the event of a sector-wide shock is sufficiently large that all firms that experience this loss go into default. Longstaff and Rajan estimate an approximately 5% loss rate on the portfolio in the event of a sector shock. Because we assume a recovery rate of 40%, the fraction of firms defaulting in the event of a sector shock is equal to $0.05/(1-.4) = 8\%$. There are 125 total firms, so a sector shock corresponds to a default of about 10 firms. Again, for simplicity, we assume that there are 25 firms in each sector (there are 125 firms, with 5 sectors); this implies a probability of 10/25, or 0.4, of firm default given a sector shock.[29]

We directly estimate the remaining parameters using moments of junior tranches in the data. To make this estimation computationally feasible, we impose some simplifying assumptions. Similar to the idiosyncratic jump size, we assume sector-wide jump size $\left(e^{Z_{S_i,t}} - 1\right)$ to be constant. While Collin-Dufresne, Goldstein, and Yang (2012) calibrate the idiosyncratic intensity for each series to match the term structure of CDX spreads (by assuming that the intensity for each series is a distinctive deterministic step function), we assume that this intensity (λ_i) is truly a constant throughout the entire sample period. These simplifying assumptions leave us only three parameters to be estimated ($Z_{S_i,t}$, w_ξ, and λ_i).

Specifically, to estimate the three parameters, we use the following four moments: (1)

[28] For simplicity we treat ϕ as a parameter rather than a process; namely the shift in ϕ corresponds to a one-time change in regime that is not anticipated by the agents.

[29] An equivalent approach is to consider a larger number of sectors, and have firm default be certain given a sector shock. This is the approach of Longstaff and Rajan (2008). The implications for the payout process are identical.

the average of pre-crisis equity tranche upfront payments, (2) the average of crisis equity tranche upfront payments, (3) the average of pre-crisis mezzanine tranche spreads, and (4) the average of crisis mezzanine tranche spreads.[30] That is, we choose parameters to minimize the sum of squared errors for these four moments. Thus we do not use data on senior tranches or on the CDX index to fit these parameters. Table 3 reports the resulting parameter values.

5.2 Implied volatilities on index options

In an earlier paper (Seo and Wachter, 2016) we establish that a special case of the model in Section 2.1 can fit average option-implied volatilities in the data. Our previous model (which has one state variable rather than two) captures the fact that implied volatilities are higher than realized volatilities and that out-of-the money (OTM) put options have higher implied volatilities than at-the-money (ATM) options. Our model explains these facts because of non-normalities arising from disasters, and high normal-times volatility arising endogenously from time-variation in the disaster probability.

Because the current model nests the previous one and is similarly calibrated, it is not surprising that this model, too, can explain average implied volatilities. Moreover, the second state variable allows it to capture time-variation in the slope of the implied volatility curve. Because we will be using the model to infer the state variables for pricing CDX tranche spreads, this property is desirable. To understand why the second factor can capture time-variation in the slope of implied volatilities, recall the process for the stock price (9). Increases in both λ_t and ξ_t increase the stock price volatility. Hence increases in these two variables

[30]Note that CDX/CDX tranche pricing requires a large number of simulations. To reduce computation time during the estimation procedure, for each date (i.e. for each pair of state variables), we construct a grid of these three variables and compute CDX/CDX tranche prices on each grid. Then, we interpolate these grid points using a 3-dimensional cubic spline. We confirm that this interpolation is very accurate because each tranche spread is a monotonic function of each parameter.

will also increase option prices, and thus option-implied volatility. Figure 1 suggests that increases in ξ_t have a larger effect on stock price volatility than increases in λ_t. On the other hand, increases in λ_t raise the immediate probability of disaster more than increases in ξ_t. We would thus expect ξ_t to have a larger effect on ATM options, and λ_t to have a larger effect on OTM options. Figure 2 shows that this is indeed the case. This figure shows 3-month implied volatilities as a function of moneyness (the strike price of the option divided by the index price) for the state variables at their median levels, and at the 20th and 80th percentiles. The level of implied volatilities is increasing in both λ_t and ξ_t. However, ξ_t affects mainly ATM implied volatilities while λ_t has a slightly greater effect for OTM implied volatilities. Increases in ξ_t that are not accompanied by increases in λ_t bring the stock-price distribution closer to log-normality, flattening the implied volatility curve.

Using the time series of option prices, we extract a time series for λ_t and ξ_t that runs from January 1996 to December 2012. While we choose these state variables to exactly fit 1-month ATM and OTM (0.85 moneyness) implied volatility, Figure 3 verifies that the fit to 3 and 6-month implied volatility is also very good. We show the time series of the extracted state variables in Panel A of Figure 4. For most of this sample, the disaster probability λ_t and its mean ξ_t lie below 5%. The state variables rise in late 1998 corresponding to the rescue of Long Term Capital Management, following the Asian financial crisis and the moratorium on payments on Russian debt. There are also increases corresponding to market declines following the NASDAQ boom in the early 2000s. The sample, however, is dominated by the financial crisis of late 2008-2009. At the Lehman default, the disaster probability rises as high as 20%. The disaster probability remains high and volatile, as compared to the prior period, until the end of the sample.

Panel B focuses in on the sample period for which we have CDX tranche data. This period captures the first hints of the crisis in early 2007, with slight increases in the level and volatility of λ_t and ξ_t. These variables become markedly higher and more volatile in 2007 and early 2008, culminating in the near-default of Bear Sterns. The CDX tranche sample ends right before the Lehman default; the last values of λ_t are higher than before, but notably lower than one month later.

Before discussing the implications of these extracted state variables for CDX pricing, we note the implications for equity prices. Table 2 already shows that the model is capable of explaining moments of equity returns. Here, we ask whether the model can capture the actual time series of stock prices. Given the option-implied values of λ_t and ξ_t, we can compute a price-dividend ratio using (8). Figure 5 shows the results, along with the price-dividend ratio from data available from Robert Shiller's webpage. The model can match the sustained level of the price-dividend ratio, and, most importantly, the time series variation after 2004, despite the fact that the state variables are derived from option prices.[31] Indeed, between 2004 and 2013, the correlation between the option-implied price-dividend ratio and the actual price-dividend ratio is 0.84, strongly suggesting these two markets share a common source of risk.

5.3 CDX/CDX tranche spreads

Given the option-implied state variables, we price CDX and CDX tranches as described in Sections 2.4 and 2.5. Recall that for the CDX and all tranches except the equity tranche, these spreads represent the annual payment per unit notional. For example, the average spread of 27 basis points for the 15-30% (3rd senior tranche) indicates that a protection buyer pays \$0.0027 per year to insure \$1.00 of value for this tranche.[32] For the equity tranche, the

[31]Not surprisingly, the disaster-risk model is not able to match the run-up in stock prices from the late 90s until around 2004. It may be that time-varying fears of a disaster will not be able to capture the extreme optimism that characterized that period.

[32]As described in Section 2.5 there are adjustments based on defaults in the portfolio.

spread is fixed at 500 basis points, so we report the upfront payment (in percentage points per unit of notional). Section 5.3.1 considers average spreads over the pre-crisis and crisis periods, while Section 5.3.2 discusses the time series.

5.3.1 Average spreads

Table 5 shows that the model can match the average spreads on the 5-year CDX and its tranches for the full sample period. More importantly, the model can match the very large shift in spreads between the pre-crisis and crisis periods. For example, on the 15-30% tranche, the average spread changes 10-fold (from 6 basis points to 69 basis points) between the samples. Our model captures both the low pre-crisis magnitude (the value is 8 basis points) and the high crisis magnitude (71 basis points). While the model generates a substantial change in the senior tranches, it correctly generates much less of a shift in the junior tranches. As we will show, pricing of the equity and mezzanine tranche has much more to do with idiosyncratic risk as opposed to market-wide risk of a rare disaster.

Table 6 shows that the model can also match spreads for the 3-year maturity. For this table, we take data values reported in Collin-Dufresne, Goldstein, and Yang (2012), who use a different data source and slightly different data construction procedure.[33] The model can match the very low average spreads prior to the crisis, and the high spreads during the crisis. The model also captures the fact that 3-year spreads are lower, but not much lower, than 5-year spreads. High 3-year spreads present a particular challenge to models that assume normally distributed risk, as discussed in Collin-Dufresne, Goldstein, and Yang (2012), and as we show in Section 4. The fact that the model can match the term structure is evidence that the correlational structure for our disaster risk process corresponds to that in the data.

Finally, Table 7 reports average spreads on the super-senior tranche. As is the case

[33] See Collin-Dufresne, Goldstein, and Yang (2012) for details. We have verified that our data and theirs are very similar for the instruments and dates where they overlap.

for Table 6, data are from Collin-Dufresne, Goldstein, and Yang (2012). The super-senior tranche is only affected if there is a loss of 30% or more on the CDX itself. Supposing the 20% recovery rate we assume for disaster periods, about 40% of the 125 firms would need to go into default for insurance on the tranche to pay off. Nowhere can rare disaster fears be seen more clearly, than in the increase in average spreads, from near zero to 30 basis points, on these tranches. Table 7 shows that the model captures the level of spreads before the crisis, the dramatic change during the crisis, and the relative spreads between the 3 and 5 year maturities. Given our focus on rare events, the ability of the model to match these average spreads is particularly encouraging. As explained in Section 4, the high crisis spreads on these products are beyond the reach of models with normally distributed risk.

5.3.2 Time-variation in spreads

We now calculate the implied time series of spreads on CDX and CDX tranches based on our state variables extracted from options data, and compare them to the historical time series. Figure 6 shows the monthly time series of five-year maturity CDX and CDX tranche spreads in the data and in the model. The blue solid line represents the data and the red dotted line represents the benchmark model. On the same figure, we also show results for the case without idiosyncratic risk, discussed further below. The top left panel presents the CDX index, and the other five panels show tranches ranging from the equity tranche to the 3rd senior tranche (15-30%). As discussed in Section 3, when the equity tranche is traded, the protection buyer makes an upfront payment to the protection seller in addition to fixed annual premium payments of 500 bps. Thus, for the equity tranche (top right panel), we report the amount of this upfront payment.

Table 5 shows that the model accurately captures the levels of the spreads on the CDX and its tranches for both the pre-crisis and crisis periods. What is new in Figure 6, and surprising, is that the timing of the increase in spreads, and the fluctuations both before and during the crisis, are accurately captured by our model. That is, the same period of low spreads and low volatility that characterized the stock and options markets in 2006 and 2007 is also apparent in the low spreads on structured finance products. The increased prices for protection, which slightly predate the collapse of Bear Sterns, appear almost simultaneously in options and CDX/CDX tranche spreads. The crisis period was one of high volatility in both markets, with the two fluctuating in tandem. The model captures the increase in ATM and OTM options over this period with an increase in both the probability of disaster, and the long-run mean of this probability. The latter increase helps the model to explain the magnitude of the increase in CDX spreads, which have a longer maturity than do options.

In fact, the only tranche that the model does not fit closely is the equity one. While the model can match the timing of the increase in the upfront payment on the equity tranche, it does not entirely capture its magnitude. Nor does it capture some of the variation in the early part of the sample, perhaps related to a credit crisis triggered by Ford and General Motors' downgrades. Nonetheless, the approximate magnitude of these spreads is well-matched, despite the fact that we have constrained idiosyncratic risk to take on the same value throughout the sample period.

We now consider the effect of changes in our assumptions on firm cash flows (12). The dashed line in Figure 6 shows the effect of setting the probability of an idiosyncratic decline in firm cash flows, λ_i, to zero. We see that the tranches respond quite differently to elimination of this risk. Specifically, the upfront payment on the equity tranche drops precipitously, to a value less than zero, indicating that the spread of 500 basis points is higher than what would be required under no upfront payment. It appears that almost all of the spreads on this tranche are due to idiosyncratic risk. As tranches increase in seniority, idiosyncratic risk has less and less of an effect. For the third senior tranche, there is no discernable effect of idiosyncratic risk at all.

Figure 7 considers the effect of eliminating sector risk (the dashed line). This limiting case is achieved by setting p_i equal to zero. Eliminating sector risk has the largest effect on tranches with intermediate levels of seniority. Without sector risk, the spreads on the mezzanine and first senior tranche are near zero for the entire sample. CDX tranche data thus clearly requires a level of commonality among firms somewhere between the firm-specific idiosyncratic risk and market-wide risk.

Figure 7 shows that the third senior tranche, with attachment point of 15% and detachment point 30%, is barely affected when the probability of a sector shock is set to zero. This confirms the intuition in Section 4 that only rare market-wide shocks have a significant effect on the pricing of this tranche. We can also see the important role of rare disasters in pricing this tranche by recalculating CDX/CDX tranche prices under a 2 percentage point increase in the probability of rare disasters at every point in time (2 percentage points is the annual mean of the disaster probability). Figure 8 shows the results. Spreads on the third senior tranche increase by about 25 basis points in the pre-crisis period, and about 50 basis points in the crisis period. The increase is substantially less than the full 2 percentage points because disaster risk is mean-reverting, and λ_t gives an instantaneous probability of disaster. Thus over the full five years of the contract, investors expect λ_t to be elevated, but by much less than 2 percentage points. The reason the effect is larger during the crisis period is because ξ_t is larger, and thus λ_t falls more slowly. Moreover, because of the square root term in (2a), higher values of λ_t (which prevail over the crisis period), are associated with greater volatility of λ_t. Because λ_t is positively skewed, this raises the probability that the high values will persist.

Finally, note that the relative effects of a change in the disaster probability on the tranches have the opposite pattern as a change in the probability of an idiosyncratic jump: the more senior the tranche, the greater the effect of a change in the probability of disaster. In particular, the top right panel of Figure 8 shows the effect on the equity tranche is negligible in proportion to the upfront payment. While an increase in the risk of a rare disaster does increase the upfront payment required by this tranche, most of this payment is insurance against idiosyncratic shocks. Only a small part represents insurance against the risk of rare disasters.

This section has shown how CDX tranche spreads and prices give insight into the correlational structure of asset values, and through the lens of our model, underlying firm earnings. Specifically, to match tranche data one needs the risk of a rare market-wide disaster, as well as fat-tailed sector-wide and idiosyncratic events. Moreover, the rare event probabilities needed to match time-series variation in CDX tranche prices are very similar to those in option prices, suggesting a rational basis for the apparently extreme pricing of these securities during the crisis.

6 Conclusion

In this paper, we build a quantitative model for spreads on the CDX and its tranches based on underlying economic fundamentals. When the model is calibrated to match the equity premium and equity volatility, and when state variables are chosen to match the time series of implied volatilities on option prices, the model can explain the level and time series of spreads both before and during the 2008–2009 financial crisis.

CDX senior tranches can be understood as extremely deep out-of-the money put options on the U.S. economy because they incur losses only when a substantial portion of large investment-grade firms default. We explain the level of spreads on these instruments by introducing a time-varying probability of economic disaster. This economic disaster causes large simultaneous declines in the consumption of representative agent, in aggregate cash flows, and in cash flows on individual firms. When agents foresee an increased probability of economic disaster, risk premia rise, asset prices decline and become more volatile, and systemic defaults become more likely. The economic disasters to which we calibrate the model are reasonable in light of what has been observed in international data over the last 100 years. The probability of such a disaster need not be high to explain senior tranche spreads: 4% is sufficient. On the other hand, we show that the spreads cannot be explained with a lognormal model for asset values; under such a model, simultaneous defaults on a significant fraction of investment-grade firms are a near-zero probability event.

Besides rare disasters, our model incorporates idiosyncratic risks and sector-wide shocks. These are important ingredients for explaining equity and mezzanine tranche prices. However, we show that the most senior tranches vary almost exclusively based on the disaster probability. Our results strongly suggest that simultaneously matching tranche prices, option prices, and properties of cash flows requires a positive probability of rare disaster that varies over time.

Appendix

Sections A–C solve for utility, aggregate market prices, and firm values in closed form up to a system of ordinary differential equations. Ratios of prices to payouts are functions of the state variables λ_t and ξ_t. Taking limits as σ_ξ approaches zero results in the model of Seo and Wachter (2016) and Wachter (2013).

A State-price density

Duffie and Skiadas (1994) show that the state-price density π_t equals

$$\pi_t = \exp\left\{\int_0^t \frac{\partial}{\partial V} f\left(C_s, V_s\right) ds\right\} \frac{\partial}{\partial C} f\left(C_t, V_t\right). \tag{A.1}$$

Our goal is to obtain an expression for the state-price density in terms of C_t, λ_t and ξ_t.

We conjecture that, in equilibrium, the continuation utility V_t equals a function J of consumption and the state variables λ_t and ξ_t such that:

$$J(C_t, \lambda_t, \xi_t) = \frac{C_t^{1-\gamma}}{1-\gamma} e^{a + b_\lambda \lambda_t + b_\xi \xi_t}. \tag{A.2}$$

For future reference, we list the derivates of J with respect to its arguments:

$$\begin{aligned}
\frac{\partial J}{\partial C} &= (1-\gamma)\frac{J}{C}, & \frac{\partial^2 J}{\partial C^2} &= -\gamma(1-\gamma)\frac{J}{C^2}, \\
\frac{\partial J}{\partial \lambda} &= b_\lambda J, & \frac{\partial^2 J}{\partial \lambda^2} &= b_\lambda^2 J, \\
\frac{\partial J}{\partial \xi} &= b_\xi J, & \frac{\partial^2 J}{\partial \xi^2} &= b_\xi^2 J.
\end{aligned} \tag{A.3}$$

Applying Ito's Lemma to $J(C_t, \lambda_t, \xi_t)$ with conjecture (A.2) and derivatives (A.3):

$$\begin{aligned}
\frac{dV_t}{V_{t^-}} = (1-\gamma)(\mu_c dt + \sigma_c dB_t) - \frac{1}{2}\gamma(1-\gamma)\sigma_c^2 dt \\
+ b_\lambda\left(\kappa_\lambda(\xi_t - \lambda_t)dt + \sigma_\lambda\sqrt{\lambda_t}dB_{\lambda,t}\right) + \frac{1}{2}b_\lambda^2\sigma_\lambda^2\lambda_t dt \\
+ b_\xi\left(\kappa_\xi(\bar{\xi} - \xi_t)dt + \sigma_\xi\sqrt{\xi_t}dB_{\xi,t}\right) + \frac{1}{2}b_\xi^2\sigma_\xi^2\xi_t dt + (e^{(1-\gamma)Z_t} - 1)dN_t.
\end{aligned}$$

Under the optimal consumption path, it must be that

$$V_t + \int_0^t f(C_s, V_s)ds = E_t\left[\int_0^\infty f(C_s, V_s)ds\right] \tag{A.4}$$

(see Duffie and Epstein (1992)). By definition,

$$\begin{aligned} f(C_t, V_t) &= \beta(1-\gamma)V_t\left(\log C_t - \frac{1}{1-\gamma}\log\left[(1-\gamma)V\right]\right) \\ &= \beta(1-\gamma)V_t \log C_t - \beta V_t \log\left[(1-\gamma)V_t\right] \\ &= \beta V_t \log\left(\frac{C_t^{1-\gamma}}{(1-\gamma)V_t}\right) \\ &= -\beta V_t(a + b_\lambda \lambda_t + b_\xi \xi_t), \end{aligned} \tag{A.5}$$

where the last equation follows from (A.2).

By the law of iterative expectations, the left-hand side of (A.4) is a martingale. Thus, the sum of the drift and the jump compensator of $(V_t + \int_0^t f(C_s, V_s)ds)$ equals zero. That is,

$$\begin{aligned} 0 = (1-\gamma)\mu_c - \frac{1}{2}\gamma(1-\gamma)\sigma_c^2 + b_\lambda\kappa_\lambda(\xi_t - \lambda_t) + \frac{1}{2}b_\lambda^2\sigma_\lambda^2\lambda_t + b_\xi\kappa_\xi(\bar{\xi} - \xi_t) + \frac{1}{2}b_\xi^2\sigma_\xi^2\xi_t \\ + \lambda_t E_\nu\left[e^{(1-\gamma)Z_t} - 1\right] - \beta(a + b_\lambda\lambda_t + b_\xi\xi_t). \end{aligned} \tag{A.6}$$

By collecting terms in (A.6), we obtain

$$\begin{aligned} 0 = \underbrace{\left[(1-\gamma)\mu_c - \frac{1}{2}\gamma(1-\gamma)\sigma_c^2 + b_\xi\kappa_\xi\bar{\xi} - \beta a\right]}_{=0} \\ + \lambda_t \underbrace{\left[-b_\lambda\kappa_\lambda + \frac{1}{2}b_\lambda^2\sigma_\lambda^2 + E_\nu\left[e^{(1-\gamma)Z_t} - 1\right] - \beta b_\lambda\right]}_{=0} \\ + \xi_t \underbrace{\left[b_\lambda\kappa_\lambda - b_\xi\kappa_\xi + \frac{1}{2}b_\xi^2\sigma_\xi^2 - \beta b_\xi\right]}_{=0}. \end{aligned} \tag{A.7}$$

Solving these equations gives us

$$a = \frac{1-\gamma}{\beta}\left(\mu_c - \frac{1}{2}\gamma\sigma_c^2\right) + \frac{b_\xi \kappa_\xi \bar{\xi}}{\beta} \tag{A.8}$$

$$b_\lambda = \frac{\kappa_\lambda + \beta}{\sigma_\lambda^2} - \sqrt{\left(\frac{\kappa_\lambda + \beta}{\sigma_\lambda^2}\right)^2 - 2\frac{E_\nu\left[e^{(1-\gamma)Z_t} - 1\right]}{\sigma_\lambda^2}} \tag{A.9}$$

$$b_\xi = \frac{\kappa_\xi + \beta}{\sigma_\xi^2} - \sqrt{\left(\frac{\kappa_\xi + \beta}{\sigma_\xi^2}\right)^2 - 2\frac{b_\lambda \kappa_\lambda}{\sigma_\xi^2}}, \tag{A.10}$$

where we have chosen the negative root based on the economic consideration that when there are no disasters, λ_t and ξ_t should not appear in the value function. Namely, for $Z_t = 0$, $b_\lambda = b_\xi = 0$. Lastly, note that these results verify the conjecture (A.2).

It follows from (A.5) that

$$\begin{aligned} \frac{\partial}{\partial C} f(C_t, V_t) &= \beta(1-\gamma)V_t C_t^{\gamma-1} C_t^{-\gamma} \\ \frac{\partial}{\partial V} f(C_t, V_t) &= \beta(1-\gamma)\left(\log C_t - \frac{1}{1-\gamma}\log\left((1-\gamma)V_t\right)\right) + \beta \end{aligned}$$

By (A.2), in equilibrium,

$$\begin{aligned} \frac{\partial}{\partial C} f(C_t, V_t) &= \beta C_t^{-\gamma} e^{a + b_\lambda \lambda_t + b_\xi \xi_t} \\ \frac{\partial}{\partial V} f(C_t, V_t) &= -\beta a - \beta - \beta b_\lambda \lambda_t - \beta b_\xi \xi_t. \end{aligned}$$

Therefore, from (A.1), it follows that the state-price density can be written as

$$\pi_t = \exp\left\{-\beta(a+1)t - \beta b_\lambda \int_0^t \lambda_s ds - \beta b_\xi \int_0^t \xi_s ds\right\} \beta C_t^{-\gamma} e^{a + b_\lambda \lambda_t + b_\xi \xi_t}. \tag{A.11}$$

B Dynamics of the aggregate market

Let $F(D_t, \lambda_t, \xi_t)$ denote the price of the dividend claim. The pricing relation implies

$$\begin{aligned} F(D_t, \lambda_t, \xi_t) &= E_t\left[\int_t^\infty \frac{\pi_s}{\pi_t} D_s ds\right] \\ &= \int_t^\infty E_t\left[\frac{\pi_s}{\pi_t} D_s\right] ds. \end{aligned}$$

Let $H(D_t, \lambda_t, \xi_t, s-t)$ denote the price of the asset that pays the aggregate dividend at time s, namely,

$$H(D_t, \lambda_t, \xi_t, s-t) = E_t\left[\frac{\pi_s}{\pi_t} D_s\right].$$

By the law of iterative expectations, it follows that $\pi_t H_t$ is a martingale:

$$\pi_t H(D_t, \lambda_t, \xi_t, s-t) = E_t[\pi_s D_s].$$

Conjecture that

$$H(D_t, \lambda_t, \xi_t, \tau) = D_t \exp\left(a_\phi(\tau) + b_{\phi\lambda}(\tau)\lambda_t + b_{\phi\xi}(\tau)\xi_t\right). \tag{B.1}$$

Applying Ito's Lemma to conjecture (B.1) implies

$$\begin{aligned}\frac{dH_t}{H_{t-}} = \Big\{ & \mu_d + b_{\phi\lambda}(\tau)\kappa_\lambda(\xi_t - \lambda_t) + \frac{1}{2}b_{\phi\lambda}(\tau)^2\sigma_\lambda^2\lambda_t + b_{\phi\xi}(\tau)\kappa_\xi(\bar{\xi} - \xi_t) + \frac{1}{2}b_{\phi\xi}(\tau)^2\sigma_\xi^2\xi_t \\ & - a'_\phi(\tau) - b'_{\phi\lambda}(\tau)\lambda_t - b'_{\phi\xi}(\tau)\xi_t \Big\} dt \\ & + \phi\sigma_c dB_t + b_{\phi\lambda}(\tau)\sigma_\lambda\sqrt{\lambda_t}dB_{\lambda,t} + b_{\phi\xi}(\tau)\sigma_\xi\sqrt{\xi_t}dB_{\xi,t} + (e^{\phi Z_t} - 1)dN_t. \end{aligned} \tag{B.2}$$

It follows from (B.2), (3), and the product rule for stochastic processes, that

$$\begin{aligned}\frac{d(\pi_t H_t)}{\pi_{t^-} H_{t^-}} = \Big\{ & -\beta - \mu_c + \gamma\sigma_c^2 - \lambda_t E_\nu\left[e^{(1-\gamma)Z_t} - 1\right] \\ & + \mu_d + b_{\phi\lambda}(\tau)\kappa_\lambda(\xi_t - \lambda_t) + \frac{1}{2}b_{\phi\lambda}(\tau)^2\sigma_\lambda^2\lambda_t \\ & + b_{\phi\xi}(\tau)\kappa_\xi(\bar{\xi} - \xi_t) + \frac{1}{2}b_{\phi\xi}(\tau)^2\sigma_\xi^2\xi_t \\ & - a'_\phi(\tau) - b'_{\phi\lambda}(\tau)\lambda_t - b'_{\phi\xi}(\tau)\xi_t \\ & - \gamma\phi\sigma_c^2 + b_\lambda b_{\phi\lambda}(\tau)\sigma_\lambda^2\lambda_t + b_\xi b_{\phi\xi}(\tau)\sigma_\xi^2\xi_t \Big\} dt \\ & + (\phi - \gamma)\sigma_c dB_t + (b_\lambda + b_{\phi\lambda}(\tau))\sigma_\lambda\sqrt{\lambda_t}dB_{\lambda,t} + (b_\xi + b_{\phi\xi}(\tau))\sigma_\xi\sqrt{\xi_t}dB_{\xi,t} \\ & + (e^{(\phi-\gamma)Z_t} - 1)dN_t. \end{aligned}$$

Since $\pi_t H_t$ is a martingale, the sum of the drift and the jump compensator of $\pi_t H_t$ equals zero. Thus:

$$
\begin{aligned}
0 = -\beta - \mu_c + \gamma\sigma_c^2 - \lambda_t E_\nu\left[e^{(1-\gamma)Z_t} - 1\right] \\
+ \mu_d + b_{\phi\lambda}(\tau)\kappa_\lambda(\xi_t - \lambda_t) + \frac{1}{2}b_{\phi\lambda}(\tau)^2\sigma_\lambda^2\lambda_t \\
+ b_{\phi\xi}(\tau)\kappa_\xi(\bar{\xi} - \xi_t) + \frac{1}{2}b_{\phi\xi}(\tau)^2\sigma_\xi^2\xi_t \\
- a'_\phi(\tau) - b'_{\phi\lambda}(\tau)\lambda_t - b'_{\phi\xi}(\tau)\xi_t \\
- \gamma\phi\sigma_c^2 + b_\lambda b_{\phi\lambda}(\tau)\sigma_\lambda^2\lambda_t + b_\xi b_{\phi\xi}(\tau)\sigma_\xi^2\xi_t + \lambda_t E_\nu\left[e^{(\phi-\gamma)Z_t} - 1\right].
\end{aligned}
\tag{B.3}
$$

Collecting terms of (B.3) results in the following equation:

$$
\begin{aligned}
0 = \underbrace{\left[-\beta - \mu_c + \gamma\sigma_c^2 + \mu_d + b_{\phi\xi}(\tau)\kappa_\xi\bar{\xi} - \gamma\phi\sigma_c^2 - a'_\phi(\tau)\right]}_{=0} \\
+ \lambda_t \underbrace{\left[-b_{\phi\lambda}(\tau)\kappa_\lambda + \frac{1}{2}b_{\phi\lambda}(\tau)^2\sigma_\lambda^2 + b_\lambda b_{\phi\lambda}(\tau)\sigma_\lambda^2 + E_\nu\left[e^{(\phi-\gamma)Z_t} - e^{(1-\gamma)Z_t}\right] - b'_{\phi\lambda}(\tau)\right]}_{=0} \\
+ \xi_t \underbrace{\left[b_{\phi\lambda}(\tau)\kappa_\lambda - b_{\phi\xi}(\tau)\kappa_\xi + \frac{1}{2}b_{\phi\xi}(\tau)^2\sigma_\xi^2 + b_\xi b_{\phi\xi}(\tau)\sigma_\xi^2 - b'_{\phi\xi}(\tau)\right]}_{=0}.
\end{aligned}
$$

It follows that

$$
\begin{aligned}
a'_\phi(\tau) &= \mu_d - \mu_c - \beta + \gamma\sigma_c^2(1-\phi) + \kappa_\xi\bar{\xi}b_{\phi\xi}(\tau) \\
b'_{\phi\lambda}(\tau) &= \frac{1}{2}\sigma_\lambda^2 b_{\phi\lambda}(\tau)^2 + (b_\lambda\sigma_\lambda^2 - \kappa_\lambda)b_{\phi\lambda}(\tau) + E_\nu\left[e^{(\phi-\gamma)Z_t} - e^{(1-\gamma)Z_t}\right] \\
b'_{\phi\xi}(\tau) &= \frac{1}{2}\sigma_\xi^2 b_{\phi\xi}(\tau)^2 + (b_\xi\sigma_\xi^2 - \kappa_\xi)b_{\phi\xi}(\tau) + \kappa_\lambda b_{\phi\lambda}(\tau).
\end{aligned}
\tag{B.4}
$$

This establishes that H satisfies the conjecture (B.1). We note that by no-arbitrage,

$$
H(D_t, \lambda_t, \xi_t, 0) = D_t.
$$

This condition provides the boundary conditions for the system of ODEs (B.4):

$$
a_\phi(0) = b_{\phi\lambda}(0) = b_{\phi\xi}(0) = 0.
$$

Finally,

$$
\begin{aligned}
F(D_t, \lambda_t, \xi_t) &= \int_t^\infty E_t\left[\frac{\pi_s}{\pi_t} D_s\right] ds \\
&= \int_t^\infty H(D_t, \lambda_t, \xi_t, s-t) ds \\
&= D_t \int_t^\infty \exp\left(a_\phi(s-t) + b_{\phi\lambda}(s-t)\lambda_t + b_{\phi\xi}(s-t)\xi_t\right) ds \\
&= D_t \int_0^\infty \exp\left(a_\phi(\tau) + b_{\phi\lambda}(\tau)\lambda_t + b_{\phi\xi}(\tau)\xi_t\right) d\tau.
\end{aligned}
$$

That is, the price-dividend ratio can be written as

$$
G(\lambda_t, \xi_t) = \int_0^\infty \exp\left(a_\phi(\tau) + b_{\phi\lambda}(\tau)\lambda_t + b_{\phi\xi}(\tau)\xi_t\right) d\tau.
$$

C Individual firm value dynamics

Let $H_i(D_{i,t}, \lambda_t, \xi_t, s-t)$ denote the time-t value of firm i's payoff at time s. That is,

$$
H_i(D_{i,t}, \lambda_t, \xi_t, s-t) = E_t\left[\frac{\pi_s}{\pi_t} D_{i,s}\right],
$$

where $D_{i,t}$ is determined by (12).[34]

We conjecture that $H_i(\cdot)$ has the following functional form:

$$
H_i(D_{i,t}, \lambda_t, \xi_t, \tau) = D_{i,t} \exp\left(a_i(\tau) + b_{i\lambda}(\tau)\lambda_t + b_{i\xi}(\tau)\xi_t\right). \tag{C.1}
$$

To verify this conjecture, we apply Ito's Lemma to the process $\pi_t H_i(D_{i,t}, \lambda_t, \xi_t, s-t)$ and derive the conditional expectation of its instantaneous change. This conditional expectation must equal zero because of (C.1), which implies that $\pi_t H_i(D_{i,t}, \lambda_t, \xi_t, s-t)$ is a martingale.

[34] In the equations that follow, we will allow Z_i and Z_{S_i} to be random variables (with independent and time-invariant distributions) rather than constants, and denote them as $Z_{i,t}$ and $Z_{S_i,t}$ respectively. $Z_{i,t}$ is assumed to be independent and identically distributed across firms, while $Z_{S_i,t}$ is independent and identically distributed across sectors.

By applying Ito's Lemma to equation (C.1), it follows that

$$
\begin{aligned}
\frac{dH_{i,t}}{H_{i,t-}} = \Bigg\{ & \mu_i + b_{i\lambda}(\tau)\kappa_\lambda(\xi_t - \lambda_t) + \frac{1}{2}b_{i\lambda}(\tau)^2\sigma_\lambda^2\lambda_t + b_{i\xi}(\tau)\kappa_\xi(\bar{\xi} - \xi_t) + \frac{1}{2}b_{i\xi}(\tau)^2\sigma_\xi^2\xi_t - a_i'(\tau) \\
& - b_{i\lambda}'(\tau)\lambda_t - b_{i\xi}'(\tau)\xi_t \Bigg\} dt + \phi_i\sigma_c dB_{c,t} + b_{i\lambda}(\tau)\sigma_\lambda\sqrt{\lambda_t}dB_{\lambda,t} + b_{i\xi}(\tau)\sigma_\xi\sqrt{\xi_t}dB_{\xi,t} \\
& + (e^{\phi_i Z_{c,t}} - 1)dN_{c,t} + I_{i,t}(e^{Z_{i,t}} - 1)dN_{i,t} + (e^{Z_{S_i,t}} - 1)dN_{S_i,t}.
\end{aligned}
$$

Note that firm i is hit by the idiosyncratic shock $dN_{i,t}$ and the sector shock $dN_{S_i,t}$. The SDE for π_t is given in (3). By applying Ito's Lemma for the product of two stochastic processes, we obtain the SDE for $\pi_t H_{i,t}$:

$$
\begin{aligned}
\frac{d(\pi_t H_{i,t})}{\pi_{t^-} H_{i,t^-}} = \Bigg\{ & -\beta - \mu_c + \gamma\sigma_c^2 - \lambda_t E\left[e^{(1-\gamma)Z_{c,t}} - 1\right] + \mu_i + b_{i\lambda}(\tau)\kappa_\lambda(\xi_t - \lambda_t) + \frac{1}{2}b_{i\lambda}(\tau)^2\sigma_\lambda^2\lambda_t \\
& + b_{i\xi}(\tau)\kappa_\xi(\bar{\xi} - \xi_t) + \frac{1}{2}b_{i\xi}(\tau)^2\sigma_\xi^2\xi_t - a_i'(\tau) - b_{i\lambda}'(\tau)\lambda_t - b_{i\xi}'(\tau)\xi_t - \gamma\phi_i\sigma_c^2 \\
& + b_\lambda b_{i\lambda}(\tau)\sigma_\lambda^2\lambda_t + b_\xi b_{i\xi}(\tau)\sigma_\xi^2\xi_t \Bigg\} dt + (\phi_i - \gamma)\sigma_c dB_{c,t} \\
& + (b_\lambda + b_{i\lambda}(\tau))\sigma_\lambda\sqrt{\lambda_t}dB_{\lambda,t} + (b_\xi + b_{i\xi}(\tau))\sigma_\xi\sqrt{\xi_t}dB_{\xi,t} \\
& + (e^{(\phi_i-\gamma)Z_{c,t}} - 1)dN_{c,t} + I_{i,t}(e^{Z_{i,t}} - 1)dN_{i,t} + (e^{Z_{S_i,t}} - 1)dN_{S_i,t}.
\end{aligned}
$$

Since $\pi_t H_t$ is a martingale, the sum of the drift and the jump compensator of $\pi_t H_t$ equals zero. This zero mean condition provides the system of ODEs for $a_i(\tau)$, $b_{i\lambda}(\tau)$, and $b_{i\xi}(\tau)$:

$$
\begin{aligned}
a_i'(\tau) &= \mu_i - \mu_c - \beta + \gamma\sigma_c^2(1 - \phi_i) + \lambda_i E\left[e^{Z_{i,t}} - 1\right] + p_i w_0 E\left[e^{Z_{S_i,t}} - 1\right] + \kappa_\xi\bar{\xi}b_{i\xi}(\tau) \\
b_{i\lambda}'(\tau) &= \frac{1}{2}\sigma_\lambda^2 b_{i\lambda}(\tau)^2 + (b_\lambda\sigma_\lambda^2 - \kappa_\lambda)b_{i\lambda}(\tau) + E\left[e^{(\phi_i-\gamma)Z_{c,t}} - e^{(1-\gamma)Z_{c,t}}\right] + p_i w_\lambda E\left[e^{Z_{S_i,t}} - 1\right] \\
b_{i\xi}'(\tau) &= \frac{1}{2}b_{i\xi}(\tau)^2\sigma_\xi^2 + (b_\xi\sigma_\xi^2 - \kappa_\xi)b_{i\xi}(\tau) + \kappa_\lambda b_{i\lambda}(\tau) + +p_i w_\xi E\left[e^{Z_{S_i,t}} - 1\right].
\end{aligned}
$$

This shows that H_i satisfies the conjecture (C.1). Furthermore, since $H_i(D_{i,t}, \lambda_t, \xi_t, 0) = D_{i,t}$, we obtain the following boundary conditions:

$$a_i(0) = b_{i\lambda}(0) = b_{i\xi}(0) = 0.$$

With the solution for the ODEs, equation (13) can be written as

$$\begin{aligned} A_i(D_{i,t},\lambda_t,\xi_t) &= \int_t^\infty H_i(D_{i,t},\lambda_t,\xi_t,s-t)\,ds \\ &= D_{i,t}\int_0^\infty \exp\left(a_i(\tau)+b_{i\lambda}(\tau)\lambda_t+b_{i\xi}(\tau)\xi_t\right)\,d\tau. \end{aligned}$$

D Computing CDX and tranche prices

Given the closed-form expressions for asset prices, the prices for the CDX and its tranches must be computed by simulation. That is, for each pair of state variables (λ_t,ξ_t), we compute the expectations that determine the protections legs (22) and (18) , and the premium legs (19) and (25) by simulating 100,000 sample paths for the 125 firms (see Appendix E for more detail on how we simulate the state variables and firm values). This exercise is complicated by the fact that, while premium payments occur at fixed (quarterly) intervals, default could occur at any time. Moreover, default is sometimes accompanied by an immediate payment, as is the case for the protection seller, and is sometimes reflected in a change in notional that matters for the premium payment at the next interval (as is the case for the premium payments on the tranches).

To reduce computation time, we compute $n_{t,s}$, $L_{t,s}$, $T^L_{j,t,s}$, and $T^R_{j,t,s}$ at quarterly intervals, which correspond to the timing of payment premium dates. Given these series, we compute the value of cash flows paid by the protection seller by assuming that default occurs at the midpoint between two premium payment dates. This follows standard practice (Mortensen, 2006), and is more accurate than simply assuming that default occurs on the premium payment date itself. Therefore

$$\begin{aligned} \mathbf{Prot}_{\mathrm{CDX}}(\lambda_t,\xi_t;T-t) &= E_t^Q\left[\int_t^T e^{-\int_t^s r_u du}dL_{t,s}\right] \\ &\simeq \sum_{m=1}^{4T} E_t^Q\left[e^{-\int_t^{t+\vartheta(m-\frac{1}{2})} r_u\,du}(L_{t,t+\vartheta m}-L_{t,t+\vartheta(m-1)})\right] \qquad \text{(D.1)} \end{aligned}$$

Because default occurs at the midpoint between two payment periods, $L_{t,t+\vartheta m} = L_{t,t+\vartheta m-\vartheta/2}$.

It will be computationally useful to write (D.1) in terms of the risk-neutral, discounted expectation of $L_{t,t+\vartheta m}$ and $L_{t,t+\vartheta(m-1)}$. To do this, we need to introduce a second approximation. Note that the riskfree rate has continuous sample paths, and ϑ is small. We therefore approximate

$$\int_{t+\vartheta(m-\frac{1}{2})}^{t+\vartheta m} r_u \, du \simeq \frac{\vartheta}{2} r_{t+\vartheta m}. \tag{D.2}$$

Combining (D.1) and (D.2), we have

$$\mathbf{Prot}_{\text{CDX}}(\lambda_t, \xi_t; T-t) \simeq$$
$$\frac{1}{4}\sum_{m=1}^{4T}\left(E_t^Q\left[e^{\frac{1}{2}\vartheta r_{t+\vartheta m}} e^{-\int_t^{t+\vartheta m} r_u du} L_{t,t+\vartheta m}\right] - E_t^Q\left[e^{-\frac{1}{2}\vartheta r_{t+\vartheta(m-1)}} e^{-\int_t^{t+\vartheta(m-1)} r_u du} L_{t,t+\vartheta(m-1)}\right]\right).$$

Recall that the premium leg equals

$$\mathbf{Prem}_{\text{CDX}}(\lambda_t, \xi_t; T-t, S) =$$
$$S E_t^Q\left[\frac{1}{4}\sum_{m=1}^{4T} e^{-\int_t^{t+\vartheta m} r_u du}(1 - n_{t,t+\vartheta m}) + \int_{t+\vartheta(m-1)}^{t+\vartheta m} e^{-\int_t^s r_u du}(s - t - \vartheta(m-1)) dn_{t,s}\right]$$
$$\simeq S \sum_{m=1}^{4T} \frac{1}{4} E_t^Q\left[e^{-\int_t^{t+\vartheta m} r_u du}(1 - n_{t,t+\vartheta m}) + e^{-\int_t^{t+(\vartheta-\frac{1}{2})m} r_u du}\left(\frac{n_{t,t+\vartheta m} - n_{t,t+\vartheta(m-1)}}{2}\right)\right].$$

where the approximation in the second equation holds if default is close to the midpoint between two premium payment dates.

Our goal is to compute the risk-neutral, discounted expectation of $n_{t,t+\vartheta m}$. Using the approximation $e^{-\int_t^{t+\vartheta(m-\frac{1}{2})} r_u du} \simeq e^{-\int_t^{t+\vartheta m} r_u du}$,

$$\begin{aligned}
&\mathbf{Prem}_{\text{CDX}}(\lambda_t, \xi_t; T-t, S) \\
&\simeq \frac{S}{4} \sum_{m=1}^{4T} E_t^Q \left[e^{-\int_t^{t+\vartheta m} r_u du} \left(1 - \frac{1}{2} n_{t,t+\vartheta m} - \frac{1}{2} n_{t,t+\vartheta(m-1)} \right) \right] \\
&= \frac{S}{4} \sum_{m=1}^{4T} \left(H_0(\lambda_t, \xi_t, \vartheta m) - \frac{1}{2} E_t^Q \left[e^{-\int_t^{t+\vartheta m} r_u du} n_{t,t+\vartheta m} \right] - \frac{1}{2} E_t^Q \left[e^{-\int_t^{t+\vartheta m} r_u du} n_{t,t+\vartheta(m-1)} \right] \right) \\
&\simeq \frac{S}{4} \sum_{m=1}^{4T} \left(H_0(\lambda_t, \xi_t, \vartheta m) - \frac{1}{2} E_t^Q \left[e^{-\int_t^{t+\vartheta m} r_u du} n_{t,t+\vartheta m} \right] \right. \\
&\qquad \left. - \frac{1}{2} E_t^Q \left[e^{-\vartheta r_{t+\vartheta(m-1)}} e^{-\int_t^{t+\vartheta(m-1)} r_u du} n_{t,t+\vartheta(m-1)} \right] \right).
\end{aligned}$$

where $H_0(\lambda_s, \xi_s, \tau)$ is the price of the default-free zero-coupon bond with maturity τ, which we derive in Appendix F. [35] we find

Like the computation for the protection leg on the CDX, our computation for the protection leg for tranche j assumes that the default occurs at the midpoint between payment periods, and uses approximation (D.2):

$$\begin{aligned}
\mathbf{Prot}_{\text{Tran},j}(\lambda_t, \xi_t, T-t) &\simeq \sum_{m=1}^{4T} E_t^Q \left[e^{-\int_t^{t+(\vartheta-\frac{1}{2})m} r_u du} (T^L_{j,t,t+\vartheta m} - T^L_{j,t,t+\vartheta(m-1)}) \right] \\
&\simeq \sum_{m=1}^{4T} \left(E_t^Q \left[e^{\frac{\vartheta}{2} r_{t+\vartheta m}} e^{-\int_t^{t+\vartheta m} r_u du} T^L_{j,t,t+\vartheta m} \right] - \right. \\
&\qquad \left. E_t^Q \left[e^{-\frac{\vartheta}{2} r_{t+\vartheta(m-1)}} e^{-\int_t^{t+\vartheta(m-1)} r_u du} T^L_{j,t,t+\vartheta(m-1)} \right] \right).
\end{aligned}$$

Recall that

$$\begin{aligned}
&\mathbf{Prem}_{\text{Tran},j}(\lambda_t, \xi_t; T-t, U, S) = \\
&\qquad U + S E_t^Q \left[\sum_{m=1}^{4T} \left(e^{-\int_t^{t+\vartheta m} r_s ds} \int_{t+\vartheta(m-1)}^{t+\vartheta m} \left(1 - T^L_{j,t,s} - T^R_{j,t,s} \right) ds \right) \right].
\end{aligned}$$

[35] In economic terms, this interest rate approximation implies that the accrued interest payment comes at the same time as the premium payment, rather than upon default. Based on the argument in Section 2.5, this implies an equivalence between the premium leg for the CDX and the premium leg for the tranches.

Under the assumption that any default occurs at the midpoint between the two payment periods, the integral above is an average:

$$\mathbf{Prem}_{\text{Tran},j}(\lambda_t, \xi_t; T-t, U, S) \simeq$$
$$U + \frac{S}{4}\sum_{m=1}^{4T} E_t^Q\left[e^{-\int_t^{t+\vartheta m} r_s ds}\frac{(1 - T^L_{j,t,t+\vartheta m} - T^R_{j,t,t+\vartheta m}) + (1 - T^L_{j,t,t+\vartheta(m-1)} - T^R_{j,t,t+\vartheta(m-1)})}{2}\right].$$

(recall $\vartheta = 1/4$). Because we want to write the premium leg in terms of risk-neutral expectations of discounted variables, we approximate

$$\begin{aligned} e^{-\vartheta r_{t+\vartheta(m-1)}} e^{-\int_t^{t+\vartheta m} r_s ds} &= e^{-\int_{t+\vartheta(m-1)}^{t+\vartheta m} r_s\, ds} e^{-\int_t^{t+\vartheta(m-1)} r_s ds} \\ &\simeq e^{-\vartheta r_{t+\vartheta(m-1)}} e^{-\int_t^{t+\vartheta(m-1)} r_s ds} \end{aligned}$$

so that

$$\begin{aligned} &\mathbf{Prem}_{\text{Tran},j}(\lambda_t, \xi_t; T-t, U, S) \simeq U + \\ &\quad \frac{S}{4}\sum_{m=1}^{4T}\Bigg(H_0(\lambda_t, \xi_t, \vartheta m) - \frac{1}{2}E_t^Q\left[e^{-\int_t^{t+\vartheta m} r_s ds} T^L_{j,t,t+\vartheta m}\right] - \frac{1}{2}E_t^Q\left[e^{-\int_t^{t+\vartheta m} r_s ds} T^R_{j,t,t+\vartheta m}\right] \\ &- \frac{1}{2}E_t^Q\left[e^{-\vartheta r_{t+\vartheta(m-1)}} e^{-\int_t^{t+\vartheta(m-1)} r_s ds} T^L_{j,t,t+\vartheta(m-1)})\right] - \frac{1}{2}E_t^Q\left[e^{-\Delta r_{t+\vartheta(m-1)}} e^{-\int_t^{t+\vartheta(m-1)} r_s ds} T^R_{j,t,t+\vartheta(m-1)}\right]\Bigg). \end{aligned}$$

Next, for any $u \in \mathbb{R}$, we define the following four expectations:

$$\begin{aligned} \text{EDR}(u, \tau, \lambda_t, \xi_t) &= E_t^Q\left[e^{u r_{t+\tau}} e^{-\int_t^{t+\tau} r_s ds} n_{t,t+\tau}\right] \\ \text{ELR}(u, \tau, \lambda_t, \xi_t) &= E_t^Q\left[e^{u r_{t+\tau}} e^{-\int_t^{t+\tau} r_s ds} L_{t,t+\tau}\right] \\ \text{ETLR}_j(u, \tau, \lambda_t, \xi_t) &= E_t^Q\left[e^{u r_{t+\tau}} e^{-\int_t^{t+\tau} r_s ds} T^L_{j,t,t+\tau}\right] \\ \text{ETRR}_j(u, \tau, \lambda_t, \xi_t) &= E_t^Q\left[e^{u r_{t+\tau}} e^{-\int_t^{t+\tau} r_s ds} T^R_{j,t,t+\tau}\right] \end{aligned} \tag{D.3}$$

We re-write the pricing formulas for the CDX index and its tranches as follows:

$$
\begin{aligned}
\mathbf{Prot}_{\text{CDX}}(\lambda_t,\xi_t;T-t) &= \sum_{m=1}^{4T}\left(\text{ELR}\left(\frac{\vartheta}{2},\vartheta m,\lambda_t,\xi_t\right)-\text{ELR}\left(-\frac{\vartheta}{2},\vartheta(m-1),\lambda_t,\xi_t\right)\right)\\
\mathbf{Prem}_{\text{CDX}}(\lambda_t,\xi_t;T-t,S) &= \frac{S}{4}\sum_{m=1}^{4T}\Bigg(H_0(\lambda_t,\xi_t,\vartheta m)-\frac{1}{2}\text{EDR}\left(0,\vartheta m,\lambda_t,\xi_t\right)\\
&\qquad -\frac{1}{2}\text{EDR}\left(\vartheta,\vartheta(m-1),\lambda_t,\xi_t\right)\Bigg)\\
\mathbf{Prot}_{\text{Tran},j}(\lambda_t,\xi_t;T-t) &= \sum_{m=1}^{4T}\left(\text{ETLR}_j\left(\frac{\vartheta}{2},\vartheta m,\lambda_t,\xi_t\right)-\text{ETLR}_j\left(-\frac{\vartheta}{2},\vartheta(m-1),\lambda_t,\xi_t\right)\right)\\
\mathbf{Prem}_{\text{Tran},j}(\lambda_t,\xi_t;T-t,U,S) &= U+\frac{S}{4}\sum_{m=1}^{4T}\Bigg(H_0(\lambda_t,\xi_t,\vartheta m)-\frac{\left[\text{ETLR}_j+\text{ETRR}_j\right]\left(0,\vartheta m,\lambda_t,\xi_t\right)}{2}\\
&\qquad -\frac{\left[\text{ETLR}_j+\text{ETRR}_j\right]\left(\vartheta,\vartheta(m-1),\lambda_t,\xi_t\right)}{2}\Bigg).
\end{aligned}
$$

To price the CDX index and its tranches, it suffices to calculate the four expectations above. Note that

$$
\text{EDR}(u,\tau,\lambda_t,\xi_t)=E_t^Q\left[e^{ur_{t+\tau}}e^{-\int_t^{t+\tau}r_s ds}n_{t,t+\tau}\right]=E_t^Q\left[e^{-\int_t^{t+\tau}r_s ds}e^{ur_{t+\tau}}n_{t,t+\tau}\right]=E_t\left[\frac{\pi_{t+\tau}}{\pi_t}e^{u\cdot r_{t+\tau}}n_{t,t+\tau}\right]
$$

and similarly for the other expectations in (D.3). It therefore suffices to calculate the physical processes for $n_{t,s}$, $L_{t,s}$, $T^R_{j,t,s}$, and $T^L_{j,t}$ at a quarterly frequency. For each value of the state variables, we do this 100,000 times to obtain the expectation. This requires 100,000 5-year simulations of the 125 firms in the index.

E Model simulation

As discussed in Appendix D, we must simulate from the model in order to price the CDX and its tranches.

The first step is to simulate a series of state variables (λ_t,ξ_t). The variable ξ_t follows the square-root process of Cox, Ingersoll, and Ross (1985), and so $\xi_{t+\Delta t}|\xi_t$ has a non-central Chi-squared distribution with $\left(\frac{4\kappa_\xi\bar{\xi}}{\sigma_\xi^2}\right)$ degrees of freedom and non-centrality parameter $\left(\frac{4\xi_t\kappa_\xi e^{-\kappa_\xi\Delta t}}{(1-e^{-\kappa_\xi\Delta t})\sigma_\xi^2}\right)$.

Over a short time interval, λ_t will be well-approximated by a CIR process. That is, we approximate the conditional distribution $\lambda_{t+\Delta t}|\lambda_t$ with a non-central Chi-squared distribution with $\left(\frac{4\kappa_\lambda\xi_t}{\sigma_\lambda^2}\right)$ degrees of freedom and non-centrality parameter $\left(\frac{4\lambda_t\kappa_\lambda e^{-\kappa_\lambda\Delta t}}{(1-e^{-\kappa_\lambda\Delta t})\sigma_\lambda^2}\right)$.[36]

Given λ_t, log consumption growth ($\log(C_{t+\Delta t}/C_t)$), and each firm's log payout growth ($\log(D_{i,t+\Delta t}/D_{i,t})$) can be drawn by discretizing the following stochastic differential equations, which follow from Ito's Lemma, applied to (1) and (12).

$$\begin{aligned} d\log C_t &= \left(\mu_c - \frac{1}{2}\sigma_c^2\right)dt + \sigma_c dB_{c,t} + Z_{c,t}N_{c,t} \\ d\log D_{i,t} &= \left(\mu_i - \frac{1}{2}\phi_i^2\sigma_c^2\right)dt + \phi_i\sigma_c dB_{c,t} + \phi_i Z_{c,t}dN_{c,t} + I_{i,t}Z_{S_i,t}dN_{S_i,t} + Z_{i,t}dN_{i,t}. \end{aligned}$$

Firm value can then be computed as

$$\begin{aligned} \frac{A_{i,t+\Delta t}}{A_{i,t}} &= \frac{D_{i,t+\Delta t}}{D_{i,t}}\frac{G_i(\lambda_{t+\Delta t},\xi_{t+\Delta t})}{G_i(\lambda_t,\xi_t)} \\ &= \exp\left[\log\left(\frac{D_{i,t+\Delta t}}{D_{i,t}}\right)\right]\frac{G_i(\lambda_{t+\Delta t},\xi_{t+\Delta t})}{G_i(\lambda_t,\xi_t)}, \end{aligned} \tag{E.1}$$

while the pricing kernel can be computed as

$$\begin{aligned} \frac{\pi_{t+\Delta t}}{\pi_t} \simeq \exp\Bigg[&\eta\Delta t - \beta b_\lambda\lambda_{t+\Delta t}\Delta t - \beta b_\xi\xi_{t+\Delta t}\Delta t \\ &-\gamma\log\left(\frac{C_{t+\Delta t}}{C_t}\right) + b_\lambda(\lambda_{t+\Delta t}-\lambda_t) + b_\xi(\xi_{t+\Delta t}-\xi_t)\Bigg]. \end{aligned} \tag{E.2}$$

Using (E.1), we can obtain a series of n_t, $L_{t,s}$, $T^L_{j,t,s}$, and TRt for all j. From these series, (E.2) and the equation for r_t, (6), we compute CDX and tranche pricing using simulations as described in Appendix D.

[36]The advantage of this approach over an Euler approximation of (2a) using a conditional normal process is that, due to the presence of ξ_t, λ_t can spend long periods of time close to zero. Thus the Euler method can lead to negative values of λ_t (which, are, strictly speaking, impossible under the model), which reduces its accuracy.

F Default-free zero-coupon bond price

Let $H_0(\lambda_t, \xi_t, s-t)$ denote the time-t price of the default-free zero-coupon bond maturing at time $s > t$. By the pricing relation,

$$H_0(\lambda_t, \xi_t, s-t) = E_t\left[\frac{\pi_s}{\pi_t}\right]. \tag{F.1}$$

By multiplying π_t on both sides of (F.1), we obtain a martingale:

$$\pi_t H_0(\lambda_t, \xi_t, s-t) = E_t\left[\pi_s\right].$$

Conjecture

$$H_0(\lambda_t, \xi_t, \tau) = \exp\left(a_0(\tau) + b_{0\lambda}(\tau)\lambda_t + b_{0\xi}(\tau)\xi_t\right). \tag{F.2}$$

By Ito's Lemma,

$$\begin{aligned}\frac{dH_{0,t}}{H_{0,t^-}} = \Bigg(& b_{0\lambda}(\tau)\kappa_\lambda(\xi_t - \lambda_t) + \frac{1}{2}b_{0\lambda}(\tau)^2\sigma_\lambda^2\lambda_t + b_{0\xi}(\tau)\kappa_\xi(\bar{\xi} - \xi_t) + \frac{1}{2}b_{0\xi}(\tau)^2\sigma_\xi^2\xi_t \\ & - a_0'(\tau) - b_{0\lambda}'(\tau)\lambda_t - b_{0\xi}'(\tau)\xi_t\Bigg)dt + b_{0\lambda}(\tau)\sigma_\lambda\sqrt{\lambda_t}dB_{\lambda,t} + b_{0\xi}(\tau)\sigma_\xi\sqrt{\xi_t}dB_{\xi,t}.\end{aligned} \tag{F.3}$$

Furthermore, we also derive the stochastic differential equation for $\pi_t H_{0,t}$ by combining equation (F.3) and (3) using Ito's Lemma:

$$\begin{aligned}\frac{d(\pi_t H_{0,t})}{\pi_{t^-}H_{0,t^-}} = \Bigg(& -\beta - \mu_c + \gamma\sigma_c^2 - \lambda_t E\left[e^{(1-\gamma)Z_{c,t}} - 1\right] \\ & + b_{0\lambda}(\tau)\kappa_\lambda(\xi_t - \lambda_t) + \frac{1}{2}b_{0\lambda}(\tau)^2\sigma_\lambda^2\lambda_t \\ & + b_{0\xi}(\tau)\kappa_\xi(\bar{\xi} - \xi_t) + \frac{1}{2}b_{0\xi}(\tau)^2\sigma_\xi^2\xi_t \\ & - a_0'(\tau) - b_{0\lambda}'(\tau)\lambda_t - b_{0\xi}'(\tau)\xi_t \\ & + b_\lambda b_{0\lambda}(\tau)\sigma_\lambda^2\lambda_t + b_\xi b_{0\xi}(\tau)\sigma_\xi^2\xi_t\Bigg)dt - \gamma\sigma_c dB_{c,t} \\ & + (b_\lambda + b_{0\lambda}(\tau))\sigma_\lambda\sqrt{\lambda_t}dB_{\lambda,t} + (b_\xi + b_{0\xi}(\tau))\sigma_\xi\sqrt{\xi_t}dB_{\xi,t} + (e^{-\gamma Z_{c,t}} - 1)dN_{c,t}.\end{aligned}$$

Since $\pi_t H_{0,t}$ is a martingale, the sum of the drift and the jump compensator of $\pi_t H_{0,t}$ equals zero. That is,

$$
\begin{aligned}
0 = -\beta - \mu_c + \gamma\sigma_c^2 - \lambda_t E\left[e^{(1-\gamma)Z_{c,t}} - 1\right] \\
+ b_{0\lambda}(\tau)\kappa_\lambda(\xi_t - \lambda_t) + \frac{1}{2}b_{0\lambda}(\tau)^2\sigma_\lambda^2\lambda_t \\
+ b_{0\xi}(\tau)\kappa_\xi(\bar{\xi} - \xi_t) + \frac{1}{2}b_{0\xi}(\tau)^2\sigma_\xi^2\xi_t \\
- a_0'(\tau) - b_{0\lambda}'(\tau)\lambda_t - b_{0\xi}'(\tau)\xi_t \\
+ b_\lambda b_{0\lambda}(\tau)\sigma_\lambda^2\lambda_t + b_\xi b_{0\xi}(\tau)\sigma_\xi^2\xi_t + \lambda_t E\left[e^{-\gamma Z_{c,t}} - 1\right].
\end{aligned}
\tag{F.4}
$$

By collecting terms of (F.4),

$$
\begin{aligned}
0 = \underbrace{\left[-\beta - \mu_c + \gamma\sigma_c^2 + b_{0\xi}(\tau)\kappa_\xi\bar{\xi} - a_0'(\tau)\right]}_{=0} \\
+ \lambda_t \underbrace{\left[-b_{0\lambda}(\tau)\kappa_\lambda + \frac{1}{2}b_{0\lambda}(\tau)^2\sigma_\lambda^2 + b_\lambda b_{0\lambda}(\tau)\sigma_\lambda^2 + E\left[e^{-\gamma Z_{c,t}} - e^{(1-\gamma)Z_{c,t}}\right] - b_{0\lambda}'(\tau)\right]}_{=0} \\
+ \xi_t \underbrace{\left[b_{0\lambda}(\tau)\kappa_\lambda - b_{0\xi}(\tau)\kappa_\xi + \frac{1}{2}b_{0\xi}(\tau)^2\sigma_\xi^2 + b_\xi b_{0\xi}(\tau)\sigma_\xi^2 - b_{0\xi}'(\tau)\right]}_{=0}.
\end{aligned}
$$

These conditions provide a system of ODEs:

$$
\begin{aligned}
a_0'(\tau) &= -\beta - \mu_c + \gamma\sigma_c^2 + b_{0\xi}(\tau)\kappa_\xi\bar{\xi} \\
b_{0\lambda}'(\tau) &= -b_{0\lambda}(\tau)\kappa_\lambda + \frac{1}{2}b_{0\lambda}(\tau)^2\sigma_\lambda^2 + b_\lambda b_{0\lambda}(\tau)\sigma_\lambda^2 + E\left[e^{-\gamma Z_{c,t}} - e^{(1-\gamma)Z_{c,t}}\right] \\
b_{0\xi}'(\tau) &= b_{0\lambda}(\tau)\kappa_\lambda - b_{0\xi}(\tau)\kappa_\xi + \frac{1}{2}b_{0\xi}(\tau)^2\sigma_\xi^2 + b_\xi b_{0\xi}(\tau)\sigma_\xi^2.
\end{aligned}
\tag{F.5}
$$

This shows that H_0 satisfies the conjecture (F.2). We can obtain the boundary conditions for the system of ODEs (F.5) because $H_0(\lambda_t, \xi_t, 0) = 1$, which is equivalent to

$$a_0(0) = b_{0\lambda}(0) = b_{0\xi}(0) = 0.$$

G Option pricing

G.1 A log-linear approximation for the price-dividend ratio

The transform analysis we use to price options requires that the log of the price-dividend ratio be linear. Fortunately, the exact price-dividend ratio we derive can be closely approximated by a log-linear function.

Let $g(\lambda, \xi) = \log G(\lambda, \xi)$. For given λ^* and ξ^*, the two-dimensional Taylor approximation implies

$$g(\lambda, \xi) \simeq g(\lambda^*, \xi^*) + \left.\frac{\partial g}{\partial \lambda}\right|_{\lambda^*, \xi^*} (\lambda - \lambda^*) + \left.\frac{\partial g}{\partial \xi}\right|_{\lambda^*, \xi^*} (\xi - \xi^*). \tag{G.1}$$

We note that

$$\begin{aligned} \left.\frac{\partial g}{\partial \lambda}\right|_{\lambda^*, \xi^*} &= \frac{1}{G(\lambda^*, \xi^*)} \left.\frac{\partial G}{\partial \lambda}\right|_{\lambda^*, \xi^*} \\ &= \frac{1}{G(\lambda^*, \xi^*)} \int_0^\infty b_{\phi\lambda}(\tau) \exp\left(a_\phi(\tau) + b_{\phi\lambda}(\tau)\lambda^* + b_{\phi\xi}(\tau)\xi^*\right) d\tau \end{aligned} \tag{G.2}$$

Similarly, we obtain

$$\begin{aligned} \left.\frac{\partial g}{\partial \xi}\right|_{\lambda^*, \xi^*} &= \frac{1}{G(\lambda^*, \xi^*)} \left.\frac{\partial G}{\partial \xi}\right|_{\lambda^*, \xi^*} \\ &= \frac{1}{G(\lambda^*, \xi^*)} \int_0^\infty b_{\phi\xi}(\tau) \exp\left(a_\phi(\tau) + b_{\phi\lambda}(\tau)\lambda^* + b_{\phi\xi}(\tau)\xi^*\right) d\tau. \end{aligned} \tag{G.3}$$

Expression (G.2) and (G.3) can be interpreted as weighted averages of the coefficients $b_{\phi\lambda}(\tau)$ and $b_{\phi\xi}(\tau)$ respectively. The average is over τ, and the weights are proportional to $\exp\left\{a_\phi(\tau) + b_{\phi\lambda}(\tau)\lambda^* + b_{\phi\xi}(\tau)\xi^*\right\}$. With this in mind, we define the notation

$$b^*_{\phi\lambda} = \frac{1}{G(\lambda^*, \xi^*)} \int_0^\infty b_{\phi\lambda}(\tau) \exp\left(a_\phi(\tau) + b_{\phi\lambda}(\tau)\lambda^* + b_{\phi\xi}(\tau)\xi^*\right) d\tau \tag{G.4}$$

$$b^*_{\phi\xi} = \frac{1}{G(\lambda^*, \xi^*)} \int_0^\infty b_{\phi\xi}(\tau) \exp\left(a_\phi(\tau) + b_{\phi\lambda}(\tau)\lambda^* + b_{\phi\xi}(\tau)\xi^*\right) d\tau, \tag{G.5}$$

and the log-linear function

$$\hat{G}(\lambda_t, \xi_t) = G(\lambda^*, \xi^*) \exp\left\{b^*_{\phi\lambda}(\lambda_t - \lambda^*) + b^*_{\phi\xi}(\xi_t - \xi^*)\right\}. \tag{G.6}$$

It follows from exponentiating both sides of (G.1) that

$$G(\lambda_t, \xi_t) \simeq \hat{G}(\lambda_t, \xi_t).$$

In our analysis, we choose λ^* and ξ^* to be $\bar{\xi}$, the stationary mean of both processes.

This log-linearization method differs from the more widely-used method of Campbell (2003), applied in continuous time by Chacko and Viceira (2005). However, in this application it is more accurate. This is not surprising, since we are able to exploit the fact that the true solution for the price-dividend ratio is known. In dynamic models with the EIS not equal to one, the solution is typically unknown.

G.2 Transform analysis

The normalized put option price is given as

$$P^n(\lambda_t, \xi_t, T-t; K^n) = E_t\left[\frac{\pi_T}{\pi_t}\left(K^n - \frac{F(D_T, \lambda_T, \xi_T)}{F(D_t, \lambda_t, \xi_t}\right)^+\right]. \tag{G.7}$$

It follows from (A.11) that

$$\frac{\pi_T}{\pi_t} = \exp\left\{\int_t^T -\beta(1 + a + b_\lambda \lambda_s + b_\xi \xi_s)\, ds - \gamma \log\left(\frac{C_T}{C_t}\right) + b_\lambda(\lambda_T - \lambda_t) + b_\xi(\xi_T - \xi_t)\right\},$$

where b_λ and b_ξ are defined by (4) and (5) respectively. It follows from $F(D_t, \lambda_t, \xi_t) = D_t G(\lambda_t, \xi_t)$, $D_t = C_t^\phi$, and (G.6) that

$$\frac{F_T}{F_t} = \exp\left\{\phi \log\left(\frac{C_T}{C_t}\right) + b^*_{\phi\lambda}(\lambda_T - \lambda_t) + b^*_{\phi\xi}(\xi_T - \xi_t)\right\},$$

where $b^*_{\phi\lambda}$ and $b^*_{\phi\xi}$ are constants defined by (G.4), and (G.5), respectively.

To use the method of Duffie, Pan, and Singleton (2000), it is helpful to write down the following stochastic process, which, under our assumptions, is well-defined for given λ_t and ξ_t:

$$X_\tau = \begin{bmatrix} \log C_{t+\tau} - \log C_t \\ \lambda_{t+\tau} \\ \xi_{t+\tau} \end{bmatrix}.$$

Note that the $\{X_\tau\}$ process is defined purely for mathematical convenience. We further define

$$d_1 = \begin{bmatrix} 0 \\ b_\lambda \\ b_\xi \end{bmatrix}, \quad d_2 = \begin{bmatrix} -\gamma \\ b_\lambda \\ b_\xi \end{bmatrix}, \quad d_3 = \begin{bmatrix} 0 \\ b^*_{\phi\lambda} \\ b^*_{\phi\xi} \end{bmatrix}, \quad d_4 = \begin{bmatrix} \phi \\ b^*_{\phi\lambda} \\ b^*_{\phi\xi} \end{bmatrix}.$$

Using this notation, (G.7) can be rewritten as

$$\begin{aligned} P^n(\lambda_t, \xi_t, T-t; K^n) = {} & K^n E_t\left[e^{-\int_0^{T-t} R(X_\tau)\, d\tau + d_2^\top X_{T-t} - d_1^\top X_0} 1_{\left\{\frac{F_T}{F_t} \leq K^n\right\}}\right] \\ & - E_t\left[e^{-\int_0^{T-t} R(X_\tau)\, d\tau + (d_2+d_4)^\top X_{T-t} - (d_1+d_3)^\top X_0} 1_{\left\{\frac{F_T}{F_t} \leq K^n\right\}}\right] \end{aligned} \quad \text{(G.8)}$$

where

$$\begin{aligned} R(X_\tau) &= \beta d_1^\top X_\tau + \beta(1+a) \\ 1_{\left\{\frac{F_T}{F_t} \leq K^n\right\}} &= 1_{\left\{d_4^\top X_{T-t} \leq \log K^n + d_3^\top X_0\right\}}. \end{aligned}$$

Since $\{X_\tau\}$ is an affine process in the sense defined by Duffie, Pan, and Singleton (2000), (G.8) characterizes the put option price in terms of expectations that can be computed using their transform analysis. Specifically, if we define

$$\mathcal{G}_{p,q}(y; X_0, T-t) \equiv E\left[e^{-\int_0^{T-t} R(X_\tau) d\tau} e^{p^\top X_{T-t}} 1_{\{q^\top X_{T-t} \leq y\}}\right], \quad \text{(G.9)}$$

then the normalized put price is expressed as

$$\begin{aligned} P^n(\lambda_t, \xi_t, T-t; K^n) = {} & e^{-d_1^\top X_0} K^n \mathcal{G}_{d_2,d_4}\left(\log K^n + d_3^\top X_0; X_0, T-t\right) \\ & - e^{-(d_1+d_3)^\top X_0} K^n \mathcal{G}_{d_2+d_4,d_4}\left(\log K^n + d_3^\top X_0; X_0, T-t\right), \end{aligned}$$

where $X_0 = [0, \lambda_t, \xi_t]$. The terms written using the function $\mathcal{G}$ can then be computed tractably using the transform analysis of Duffie et al: this analysis requires only the solution of a system of ordinary differential equations and a one-dimensional numerical integration.

References

Abel, Andrew, 1990, Asset prices under habit formation and catching up with the Joneses, *American Economic Review Papers and Proceedings* 80, 38–42.

Bansal, Ravi, and Amir Yaron, 2004, Risks for the long-run: A potential resolution of asset pricing puzzles, *Journal of Finance* 59, 1481–1509.

Barro, Robert J., 2006, Rare disasters and asset markets in the twentieth century, *Quarterly Journal of Economics* 121, 823–866.

Barro, Robert J., and José F. Ursúa, 2008, Macroeconomic crises since 1870, *Brookings Papers on Economic Activity* no. 1, 255–350.

Black, Fischer, and John C Cox, 1976, Valuing corporate securities: Some effects of bond indenture provisions, *The Journal of Finance* 31, 351–367.

Black, Fischer, and Myron Scholes, 1973, The Pricing of Options and Corporate Liabilities, *Journal of Political Economy* 81, 637–654.

Borovička, Jaroslav, Lars Peter Hansen, Mark Hendricks, and José A. Scheinkman, 2011, Risk-price dynamics, *Journal of Financial Econometrics* 9, 3–65.

Campbell, John Y., 2003, Consumption-based asset pricing, in G. Constantinides, M. Harris, and R. Stulz, eds.: *Handbook of the Economics of Finance, vol. 1b* (Elsevier Science, North-Holland).

Campbell, John Y., and John H. Cochrane, 1999, By force of habit: A consumption-based explanation of aggregate stock market behavior, *Journal of Political Economy* 107, 205–251.

Campbell, John Y., and Glen B. Taksler, 2003, Equity Volatility and Corporate Bond Yields, *The Journal of Finance* 58, 2321–2349.

Chacko, George, and Luis Viceira, 2005, Dynamic consumption and portfolio choice with stochastic volatility in incomplete markets, *Review of Financial Studies* 18, 1369–1402.

Christoffersen, Peter, Du Du, and Redouane Elkamhi, 2016, Rare Disasters, Credit, and Option Market Puzzles, *Management Science*.

Collin-Dufresne, Pierre, Robert S Goldstein, and Fan Yang, 2012, On the Relative Pricing of Long-Maturity Index Options and Collateralized Debt Obligations, *The Journal of Finance* 67, 1983–2014.

Coval, Joshua D, Jakub W Jurek, and Erik Stafford, 2009, Economic catastrophe bonds, *The American Economic Review* pp. 628–666.

Cox, John C., Jonathan C. Ingersoll, and Stephen A. Ross, 1985, A theory of the term structure of interest rates, *Econometrica* 53, 385–408.

Culp, Christopher L., Yoshio Nozawa, and Pietro Veronesi, 2014, Option-based credit spreads, NBER Working Paper No. 20776.

Duffie, Darrell, 2001, *Dynamic Asset Pricing Theory.* (Princeton University Press Princeton, NJ) 3 edn.

Duffie, Darrell, and Larry G Epstein, 1992, Asset pricing with stochastic differential utility, *Review of Financial Studies* 5, 411–436.

Duffie, Darrell, and Nicolae Garleanu, 2001, Risk and valuation of collateralized debt obligations, *Financial Analysts Journal* 57, 41–59.

Duffie, Darrell, Jun Pan, and Kenneth Singleton, 2000, Transform analysis and asset pricing for affine jump-diffusions, *Econometrica* 68, 1343–1376.

Duffie, Darrell, and Kenneth J. Singleton, 1997, An Econometric Model of the Term Structure of Interest-Rate Swap Yields, *The Journal of Finance* 52, 1287–1321.

Duffie, Darrell, and Costis Skiadas, 1994, Continuous-time asset pricing: A utility gradient approach, *Journal of Mathematical Economics* 23, 107–132.

Epstein, Larry, and Stan Zin, 1989, Substitution, risk aversion and the temporal behavior of consumption and asset returns: A theoretical framework, *Econometrica* 57, 937–969.

Fama, Eugene F., and Kenneth R. French, 2001, Disappearing dividends: Changing firm characteristics or lower propensity to pay?, *Journal of Financial Economics* 60, 3–43.

Gabaix, Xavier, 2012, An exactly solved framework for ten puzzles in macro-finance, *Quarterly Journal of Economics* 127, 645–700.

Gorton, Gary, and Andrew Metrick, 2012, Securitized banking and the run on repo, *Journal of Financial Economics* 104, 425–451.

Gourio, François, 2013, Credit risk and disaster risk, *American Economic Journal: Macroeconomics* 5, 1–34.

Kelly, Bryan T., Gerardo Manzo, and Diogo Palhares, 2016, Credit-Implied Volatility, Working paper, University of Chicago.

Lemmon, Michael L., Michael R. Roberts, and Jaime F. Zender, 2008, Back to the beginning: persistence and the cross-section of corporate capital structure, *The Journal of Finance* 63, 1575–1608.

Lettau, Martin, and Jessica A. Wachter, 2007, Why is long-horizon equity less risky? A duration-based explanation of the value premium, *Journal of Finance* 62, 55–92.

Longstaff, Francis A., and Monika Piazzesi, 2004, Corporate earnings and the equity premium, *Journal of Financial Economics* 74, 401–421.

Longstaff, Francis A, and Arvind Rajan, 2008, An empirical analysis of the pricing of collateralized debt obligations, *The Journal of Finance* 63, 529–563.

Moreira, Alan, and Alexi Savov, 2016, The Macroeconomics of Shadow Banking, Working paper, New York University and Yale University.

Mortensen, Allan, 2006, Semi-Analytical Valuation of Basket Credit Derivatives in Intensity-Based Models., *Journal of Derivatives* 13, 8–26.

Reinhart, Carmen M., and Kenneth S. Rogoff, 2009, *This time is different: eight centuries of financial folly.* (Princeton University Press Princeton, NJ).

Rietz, Thomas A., 1988, The equity risk premium: A solution, *Journal of Monetary Economics* 22, 117–131.

Salmon, Felix, 2009, Recipe for Disaster: The Formula that Killed Wall Street, *Wired Magazine.*

Seo, Sang Byung, 2014, Correlated defaults and economic catastrophes: Link between CDS market and asset returns, Working paper, University of Houston.

Seo, Sang Byung, and Jessica A. Wachter, 2016, Option prices in a model with stochastic disaster risk, NBER Working Paper #19611.

Tsai, Jerry, and Jessica A. Wachter, 2015, Disaster Risk and Its Implications for Asset Pricing, *Annual Review of Financial Economics* 7, 219–252.

Wachter, Jessica A., 2013, Can time-varying risk of rare disasters explain aggregate stock market volatility?, *The Journal of Finance* 68, 987–1035.

Weil, Philippe, 1990, Nonexpected utility in macroeconomics, *Quarterly Journal of Economics* 105, 29–42.

Zhou, Chunsheng, 2001, The term structure of credit spreads with jump risk, *Journal of Banking & Finance* 25, 2015–2040.

Figure 1: Solution for the price-dividend ratio

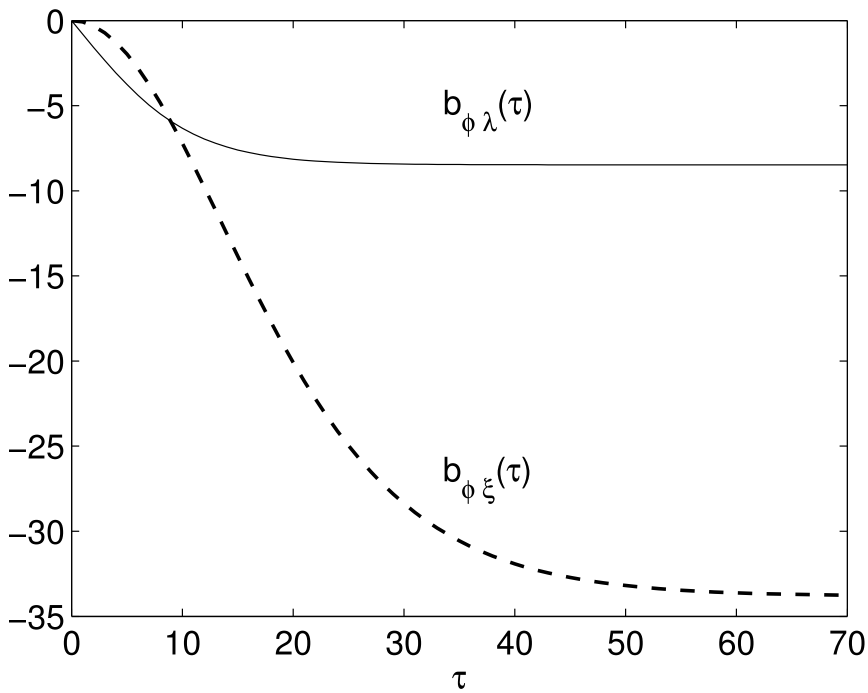

Notes: The functions $b_{\phi\lambda}(\tau)$ and $b_{\phi\xi}(\tau)$, which determine the sensitivity of the aggregate market to changes in the disaster probability λ_t and to its time-varying mean ξ_t. That is, the price-dividend ratio on the aggregate market is given by $G(\lambda_t, \xi_t) = \int_0^\infty \exp\left(a_\phi(\tau) + b_{\phi\lambda}(\tau)\lambda_t + b_{\phi\xi}(\tau)\xi_t\right) d\tau$. The horizon τ is in years.

Figure 2: Implied volatilities in the model

Panel A: Varying λ_t

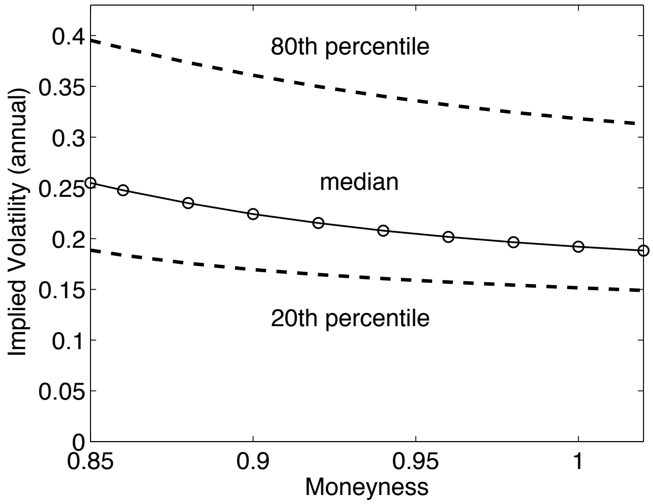

Panel B: Varying ξ_t

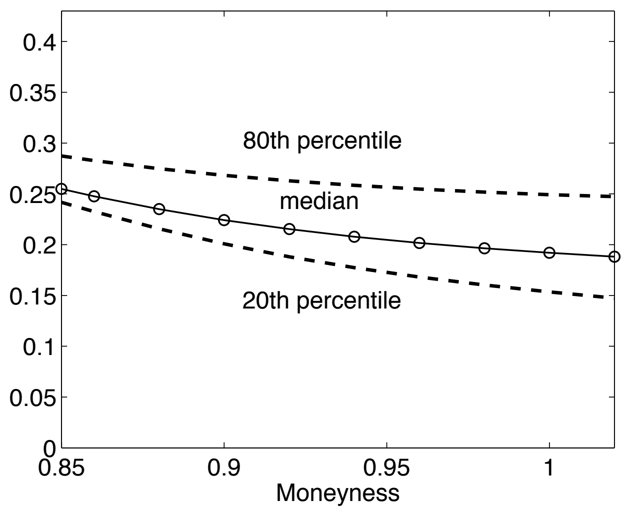

Notes: Implied volatilities for 3-month put options on the equity index, shown as a function of moneyness (the strike price divided by the index price), as calculated in the model. The figures show the effects of varying the state variables λ_t (the disaster probability) and ξ_t (the value to which λ_t reverts). Panel A sets ξ_t equal to its median value and varies λ_t, while Panel B sets λ_t equal to its median value and varies ξ_t.

Figure 3: Time series of option-implied volatilities

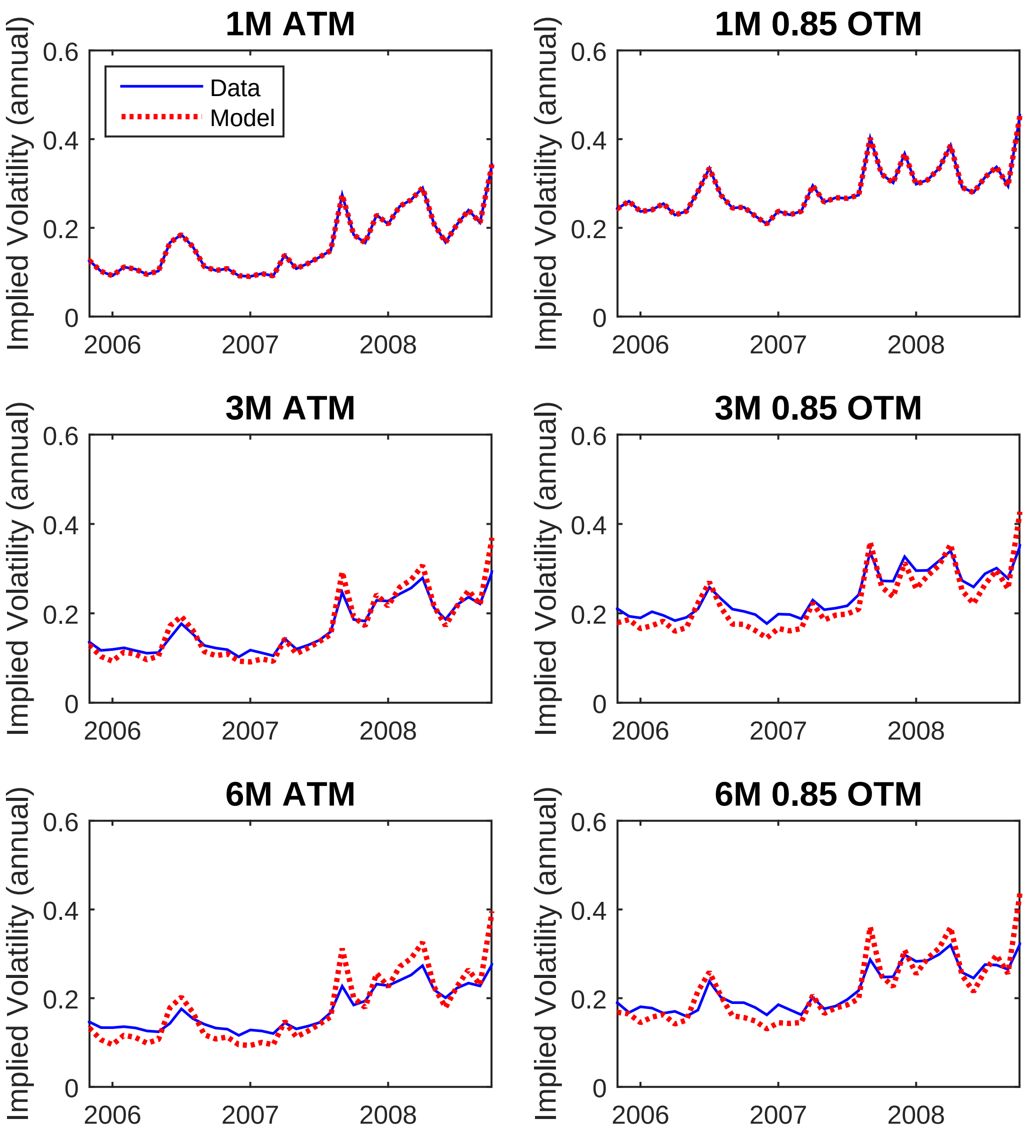

Notes: Monthly time series of option-implied volatilities in the data (blue solid lines) and in the model (red dotted lines). Results are shown for 1, 3, and 6-month options. State variables are computed to match the 1-month ATM and 0.85 OTM implied volatilities exactly.

Figure 4: Option-implied state variables

Panel A: Option sample period

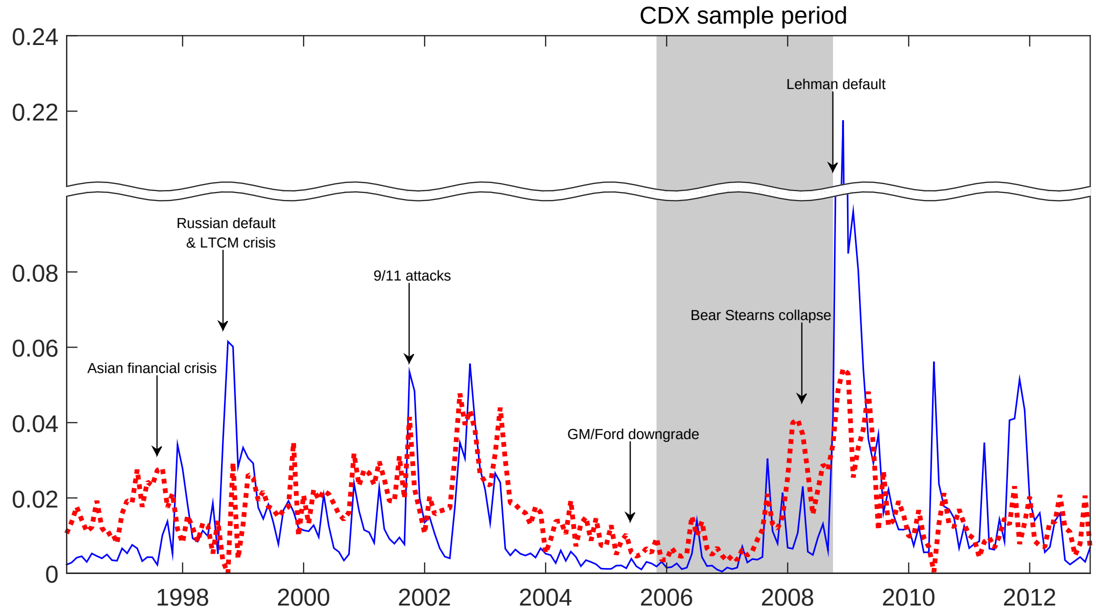

Panel B: CDX tranche sample period

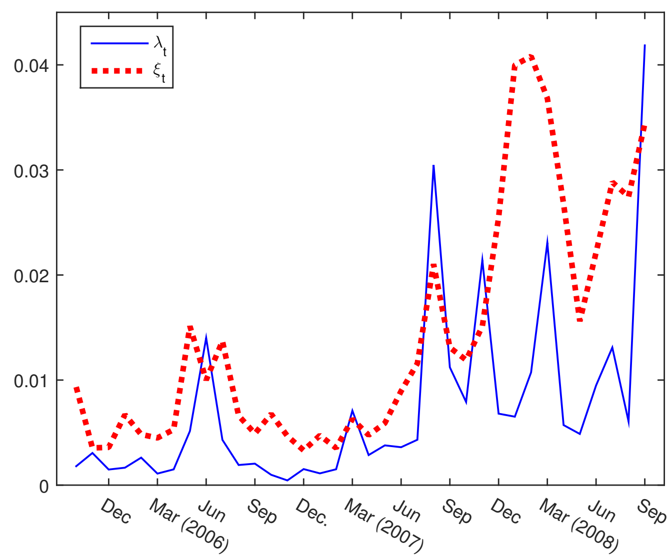

Notes: Monthly time series of the state variables λ_t (annual disaster probability) and ξ_t (long-run mean of λ_t) extracted from option prices. At each time point, the state variables are chosen to match implied volatilities of 1-month ATM and OTM (moneyness of 0.85) index put options in the data. Panel A shows these time series for the January 1996–December 2012 period over which option data are available. The shaded area represents the period over with CDX tranche data are available. Panel B zooms into the CDX sample period.

Figure 5: The price-dividend ratio in the data and in the model

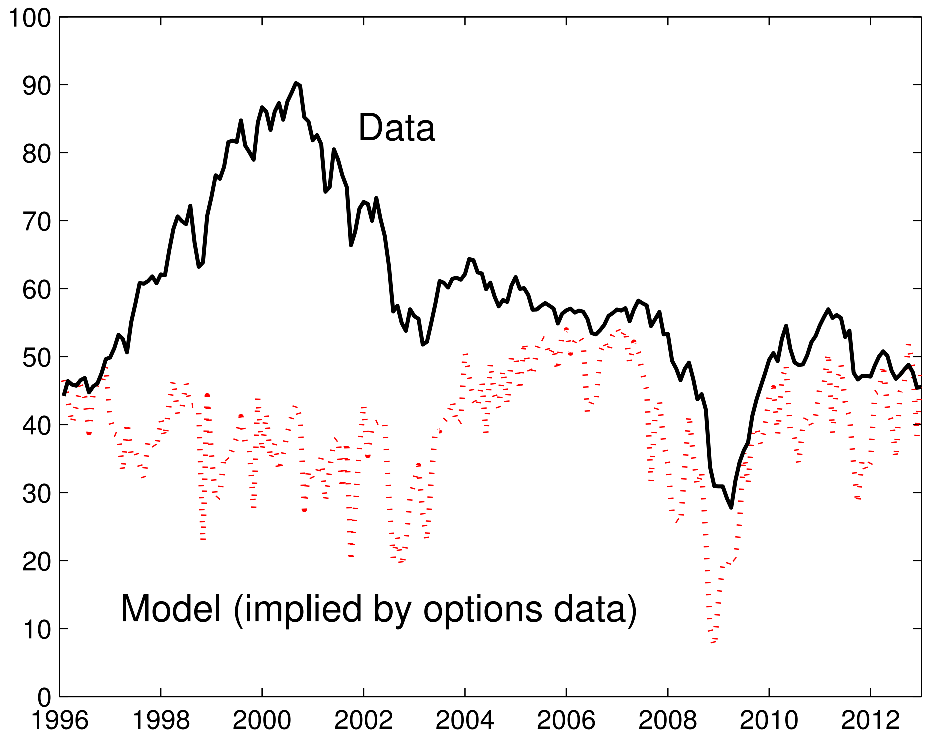

Notes: The solid line shows the time series of the price-dividend ratio on the data. The red line shows the price-dividend ratio implied by the model for state prices chosen to fit the one-month ATM and OTM (0.85 moneyness) put options.

Figure 6: Time series of CDX/CDX tranche spreads in the benchmark model and in a case without idiosyncratic risk

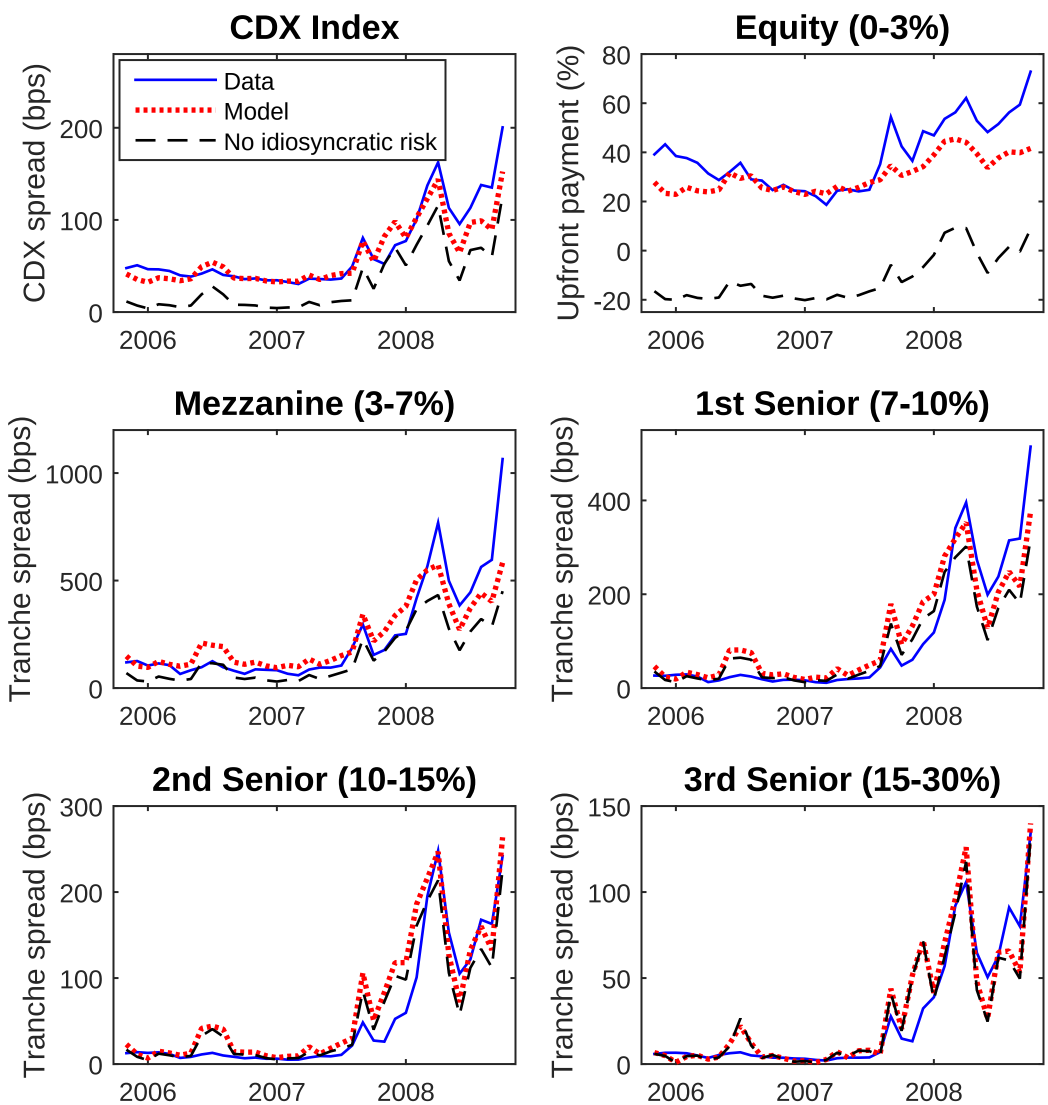

Notes: Monthly time series of 5-year CDX and CDX tranche spreads in the data (blue solid lines), in the benchmark model (red dotted lines), and in a calibration without idiosyncratic risk (black dashed lines). Spreads are annual and reported in terms of basis points (bps) per unit of notional. For the equity tranche, we report the upfront payment assuming the spread is fixed at 500 bps. We compute model values from state variables fit to the time series of 1-month ATM and 0.85 OTM implied volatilities. The version without idiosyncratic risk is obtained from the benchmark by setting the probability of an idiosyncratic shock (λ_i) to zero.

Figure 7: Time series of CDX/CDX tranche spreads in the benchmark model and in a case without sector shocks

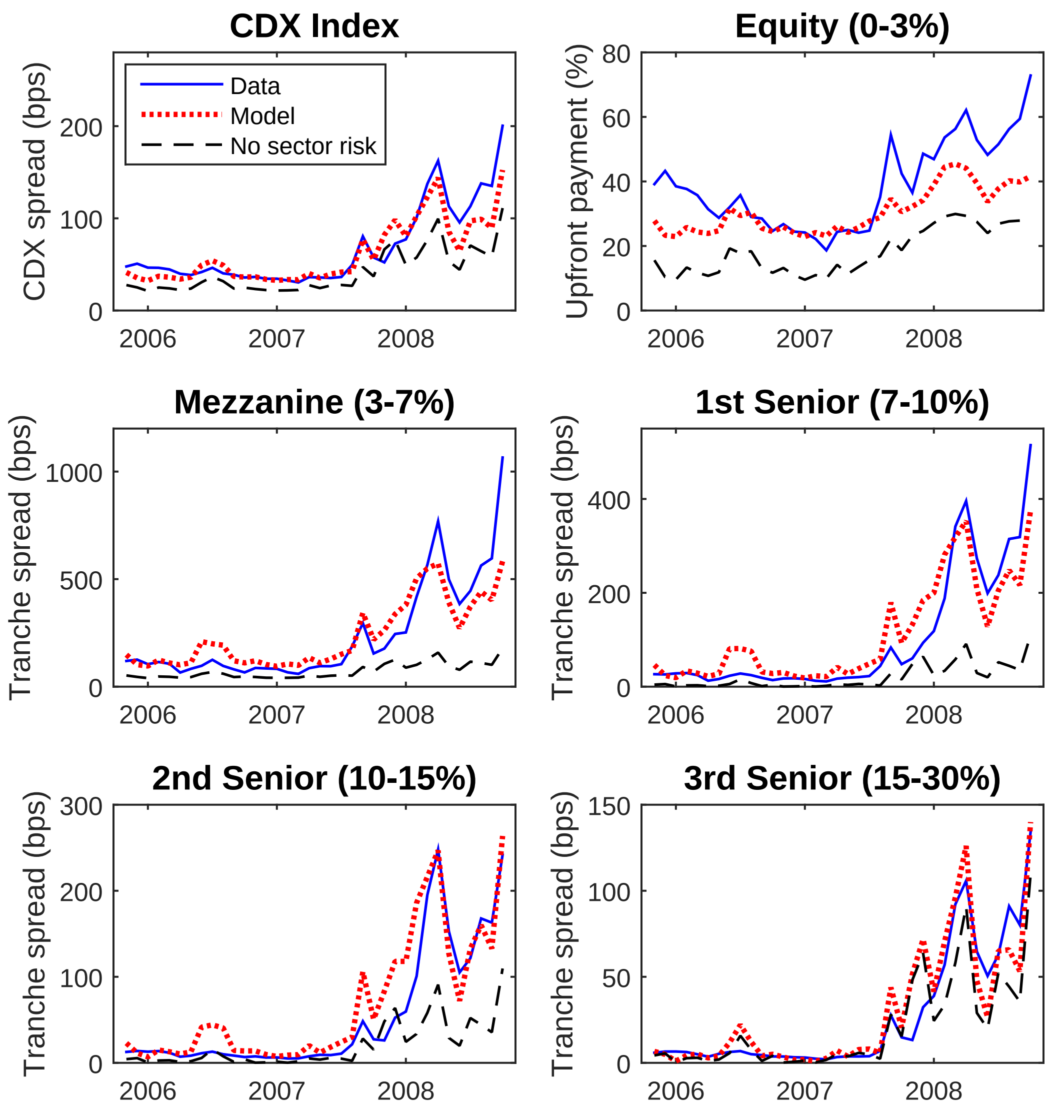

Notes: Monthly time series of 5-year CDX and CDX tranche spreads in the data (blue solid lines), in the benchmark model (red dotted lines), and in a calibration without sector shocks (black dashed lines). Spreads are annual and reported in terms of basis points (bps) per unit of notional. For the equity tranche, we report the upfront payment assuming the spread is fixed at 500 bps. We compute model values from state variables fit to the time series of 1-month ATM and 0.85 OTM implied volatilities. The version without sector risk is obtained by setting the probability that a firm will suffer a sector shock (p_i) to zero.

Figure 8: Time series of CDX/CDX tranche spreads in the benchmark model and in a case with higher disaster risk

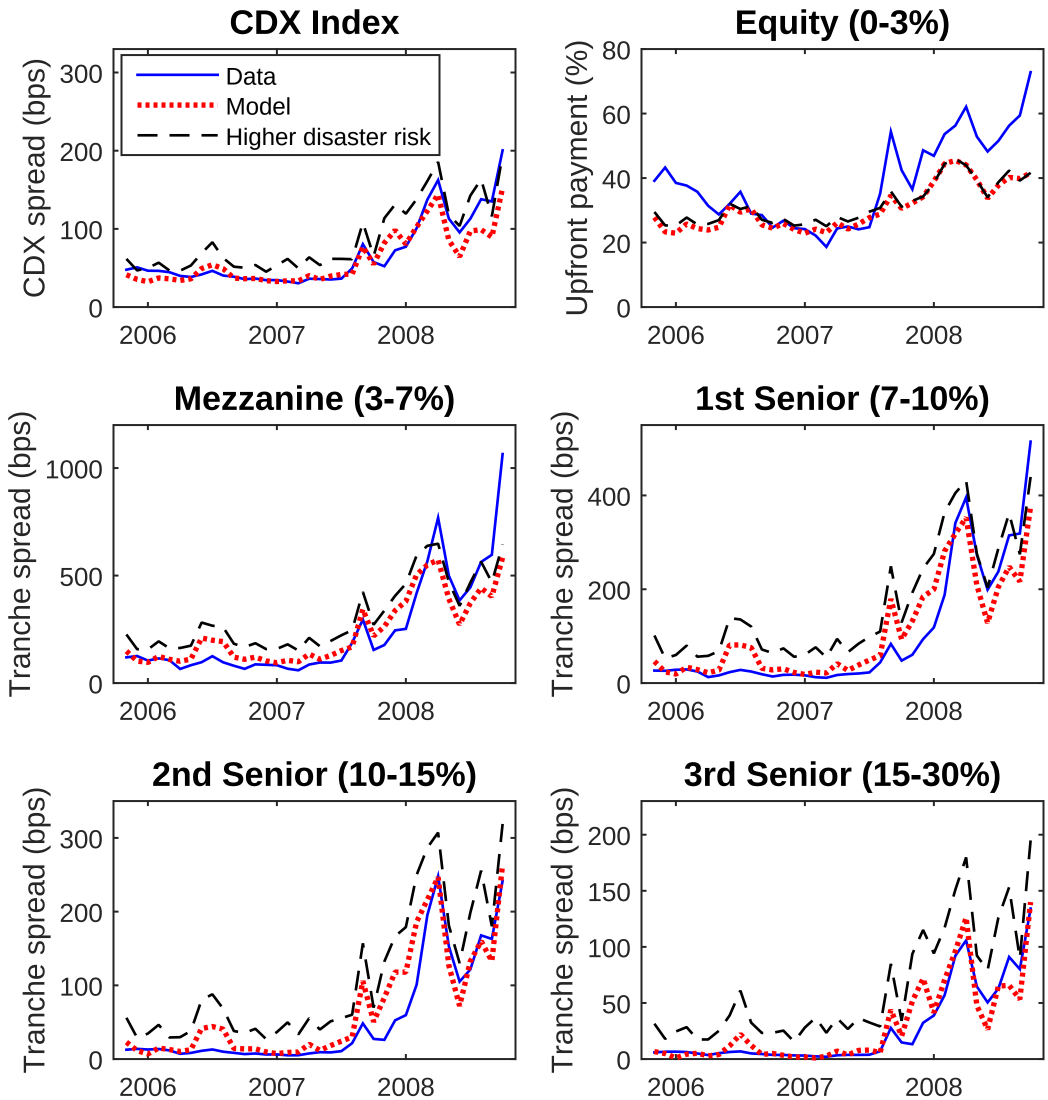

Notes: Monthly time series of 5-year CDX and CDX tranche spreads in the data (blue solid lines), in the benchmark model (red dotted lines), and in a calibration with higher disaster risk. Spreads are reported in terms of basis points (bps). For the equity tranche, we report the upfront payment because the spread is fixed at 500 bps. The benchmark model values are computed using the option-implied state variables fit to the time series of 1-month ATM and 0.85 OTM implied volatilities. The calibration with higher disaster risk is obtained by uniformly increasing the disaster probability λ_t by 2 percentage points.

Table 1: Properties of aggregate cash flows and utility

Panel A: Parameters for utility, consumption, and dividends	
Relative risk aversion γ	3
EIS ψ	1
Rate of time preference β	0.012
Average growth in consumption (normal times) μ_c	0.0252
Volatility of consumption growth (normal times) σ_c	0.020
Leverage ϕ	2.6
Panel B: Parameters for the disaster probability process	
Mean reversion κ_λ	0.20
Volatility parameter σ_λ	0.1576
Mean reversion κ_ξ	0.10
Volatility parameter σ_ξ	0.0606
Mean $\bar{\xi}$	0.02
Panel C: Population statistics of the disaster probability λ_t	
Median	0.0037
Standard deviation	0.0386
Monthly autocorrelation	0.9858

Notes: Panel A shows parameters for normal-times consumption and dividend processes, and for the preferences of the representative agent. Panels B shows the parameter values for λ and ξ processes:

$$
\begin{aligned}
d\lambda_t &= \kappa_\lambda(\xi_t - \lambda_t)dt + \sigma_\lambda\sqrt{\lambda_t}dB_{\lambda,t} \\
d\xi_t &= \kappa_\xi(\bar{\xi} - \xi_t)dt + \sigma_\xi\sqrt{\xi_t}dB_{\xi,t}
\end{aligned}
$$

Note that $\bar{\xi}$ is the average level of the probability of a disaster. Panel C shows population statistics for the disaster probability λ_t. Except for the autocorrelation of λ_t, values are in annual terms.

Table 2: Moments for the government bill rate and the market return in the data and the model

		No-Disaster Simulations			All Simulations			
	Data	0.05	0.50	0.95	0.05	0.50	0.95	Population
$E[R^b]$	1.25	1.68	2.96	3.46	-0.47	2.41	3.37	2.02
$\sigma(R^b)$	2.75	0.34	1.07	2.71	0.48	2.06	7.14	3.69
$E[R^m - R^b]$	7.25	5.40	8.01	12.36	5.30	8.49	14.25	9.00
$\sigma(R^m)$	17.8	13.24	19.26	27.91	14.59	22.59	34.38	24.13
Sharpe Ratio	0.41	0.32	0.42	0.55	0.26	0.39	0.53	0.37
$\exp(E[p-d])$	32.5	28.96	40.63	48.88	22.93	36.95	47.41	35.36
$\sigma(p-d)$	0.43	0.15	0.27	0.47	0.17	0.33	0.59	0.43
AR1$(p-d)$	0.92	0.59	0.79	0.91	0.62	0.82	0.92	0.90

Notes: Data moments are calculated using annual data from 1947 to 2010. Population moments are calculated from simulating data from the model at a monthly frequency for 600,000 years and then aggregating monthly growth rates to an annual frequency. We also simulate 100,000 60-year samples and report the 5th, 50th and 95th percentile for each statistic, both from the full set of simulations and for the subset of samples for which no disasters occur. R^b denotes the government bill return, R^m denotes the return on the aggregate market and $p - d$ denotes the log price-dividend ratio.

Table 3: Parameter values for an individual firm

Default boundary A_B	19.2%
Recovery rate R_t (normal times)	40%
Recovery rate R_t (disaster times)	20%
Aggregate risk loading ϕ_i (pre-crisis)	1.3
Aggregate risk loading ϕ_i (crisis)	1.6
Idiosyncratic jump size $\left(e^{Z_i} - 1\right)$	-80.0%
Bernoulli parameter for sector shocks $P(I_{it} = 1) = p_i$	0.4
Sector-wide jump size $\left(e^{Z_{S_i}} - 1\right)$	-71.1%
Coefficient for sector-wide jump intensity w_ξ	1.765
Idiosyncratic jump intensity λ_i	0.0093

Notes: This table reports the parameters for the payout process on an individual firm. Note that default boundary A_B is calculated as the firm's leverage ratio (0.32) multilied by 60% following Collin-Dufresne, Goldstein, and Yang (2012). Idiosyncratic jump intensity λ_i is in annual terms.

Table 4: Default probabilities under a lognormal model for prices

	Drift					
	-6%	-4%	-2%	-6%	-4%	-2%
Volatility	5-year probability			3-year probability		
14%	0.0027	0.0007	0.0001	<.0001	<.0001	<.0001
16%	0.0269	0.0093	0.0030	<.0001	<.0001	<.0001
18%	0.1335	0.0580	0.0238	0.0005	0.0002	0.0001
20%	0.4264	0.2182	0.1068	0.0041	0.0020	0.0009

Notes: The probability of hitting the default boundary for a firm subject to diffusive risk only, for volatility ranging from 14% to 20% per annum and drift ranging from -6% to -2% per annum. Volatility is for firm value (equity plus debt); the drift is under the risk-neutral measure and so should be interpreted as the riskfree rate minus the ratio of total payout to firm value. Values are in percentages (meaning that 1 denotes a 1% probability and 100 denotes certainty). The values on the left give the probability of default within 5 years. The values on the right give the probability of default within 3 years.

Table 5: Average CDX and CDX tranche spreads (5-year maturity)

	Upfront (%)	Annual Spread (bps)				
	0-3%	3-7%	7-10%	10-15%	15-30%	CDX
Panel A: Pre-crisis (Oct. 2005 – Sep. 2007)						
Data	31	108	25	12	6	42
Model	26	142	46	23	8	41
Panel B: Crisis (Oct. 2007 – Sep. 2008)						
Data	54	498	255	136	69	116
Model	39	422	238	155	71	102
Panel C: Full sample (Oct. 2005 – Sep. 2008)						
Data	39	238	102	54	27	67
Model	31	235	110	67	29	61

Notes: Historical and model-implied average 5-year CDX and CDX tranche spreads in basis points (bps) per year. For the equity tranche (0-3%), the spread is fixed 500 bps, so we report the upfront payment. Other tranches (3-7%, 7-10%, 10-15%, 15-30%) and the CDX itself have no upfront payments. Data are monthly, from October 2005 to September 2008, corresponding to CDX series 5 to 10. We divide the data into two sub-samples: pre-crisis and crisis. The The pre-crisis sample is from October 2005 to September 2007 (CDX5 to CDX8). The crisis sample is from October 2007 to September 2008 (CDX9 to CDX10). We compute model values from state variables fit to the time series of 1-month ATM and 0.85 OTM implied volatilities on equity index options.

Table 6: Average CDX and CDX tranche spreads (3-year maturity)

	Upfront (%)	Annual Spread (bps)				
	0-3%	3-7%	7-10%	10-15%	15-30%	CDX
Panel A: Pre-crisis (Sep. 2004 – Sep. 2007)						
Data	11	20	8	3	2	27
Model	14	42	13	6	3	32
Panel B: Crisis (Oct. 2007 – Sep. 2008)						
Data	43	364	168	87	48	–
Model	24	243	130	80	46	82
Panel C: Full sample (Sep. 2004 – Sep. 2008)						
Data	21	127	58	29	16	–
Model	17	91	42	24	13	44

Notes: Historical and model-implied average 3-year CDX and CDX tranche spreads in basis points (bps) per year. For the equity tranche (0-3%), the spread is fixed 500 bps, so we report the upfront payment. Other tranches (3-7%, 7-10%, 10-15%, 15-30%) and the CDX itself have no upfront payments. Data, from Collin-Dufresne, Goldstein, and Yang (2012), are monthly from September 2004 to September 2008. The pre-crisis sample is from September 2004 to September 2007 (CDX3 to CDX8). The crisis sample is from October 2007 to September 2008 (CDX9 and 10). We compute model values from state variables fit to the time series of 1-month ATM and 0.85 OTM implied volatilities on equity index options.

Table 7: Average super-senior tranche spreads

	Super senior tranche spread (annual bps)			
	Pre-crisis (Sep. 2004 – Sep. 2007)		Crisis (Oct. 2007 – Sep. 2008)	
	3-year	5-year	3-year	5-year
Data	1	4	23	35
Model	2	3	28	34

Notes: Historical and model-implied average 3 and 5-year super-senior tranche spreads in basis points per year. The super-senior tranche has a 30% attachment point and a 100% detachment point. Data, from Collin-Dufresne, Goldstein, and Yang (2012), are monthly from September 2004 to September 2008. The pre-crisis sample is from September 2004 to September 2007 (CDX3 to CDX8). The crisis sample is from October 2007 to September 2008 (CDX9 and 10). We compute model values from state variables fit to the time series of 1-month ATM and 0.85 OTM implied volatilities on equity index options.