The Journal of FINANCE
The Journal of THE AMERICAN FINANCE ASSOCIATION

THE JOURNAL OF FINANCE • VOL. LXXIV, NO. 4 • AUGUST 2019

Time-Varying Asset Volatility and the Credit Spread Puzzle

DU DU, REDOUANE ELKAMHI, and JAN ERICSSON*

ABSTRACT

Most extant structural credit risk models underestimate credit spreads—a shortcoming known as the credit spread puzzle. We consider a model with priced stochastic asset risk that is able to fit medium- to long-term spreads. The model, augmented by jumps to help explain short-term spreads, is estimated on firm-level data and identifies significant asset variance risk premia. An important feature of the model is the significant time variation in risk premia induced by the uncertainty about asset risk. Various extensions are considered, among them optimal leverage and endogenous default.

STRUCTURAL CREDIT RISK MODELS HAVE MET WITH SIGNIFICANT difficulties in academic research. First, attempts to empirically implement models on individual corporate bond prices have failed.[1] Second, efforts to calibrate models to observable moments including historical default rates and Sharpe ratios have been unable to match average credit spreads levels (the credit spread puzzle; Huang and Huang (2012)). Finally, models have been unable to jointly explain dynamics of credit spreads and equity volatilities (Huang and Zhou (2008)).

An important recent insight is that model improvements are likely to come from modeling risk premia rather than default probabilities (Chen,

*Du Du is with the City University Hong Kong. Redouane Elkamhi is with the Rotman School of Management at the University of Toronto. Jan Ericsson is with Desautels Faculty of Management at McGill University. Min Jiang contributed to an earlier version of this paper. Ericsson has benefited from financial support from a Desmarais Faculty Scholarship, from the Institut de Finance Mathématique de Montréal, as well as SSHRC. Du has benefited from financial support by the National Natural Science Foundation of China (No. 71720107002). The paper has benefited from comments by seminar participants at Case Western Reserve University, HEC Montreal, Hong Kong University of Science and Technology, University of Houston, University of Iowa, Konstanz University, National University of Singapore, and University of Toronto and conference participants at the 9th International Paris Finance Meeting 2011, ITAM 2012, Tremblant Risk Management 2012, Risk Management Conference NUS 2011, SKK, IFSID 2012, Montreal, and Affi 2016, Liège. Special thanks to David Bates, Peter Christoffersen, Sudipto Dasgupta, Jin Chuan Duan, Adlai Fisher (discussant), Santiago García Verdú, Kris Jacobs, Chanik Jo, Aytek Malkhozov, Franck Moraux (discussant), Raunaq Pungaliya, Tom Rietz, Sheikh Sadik, Ashish Tiwari, Anand Vijh, Hao Wang, and Hao Zhou (discussant). We thank Evan Zhou for exceptional research assistance. The authors have read the *Journal of Finance*'s disclosure policy and have no conflicts of interest to declare.

[1] See Jones, Mason, and Rosenfeld (1984) and Eom, Helwege, and Huang (2004).

DOI: 10.1111/jofi.12765

Collin-Dufresne, and Goldstein (2009, CCG)). Building on this insight, we develop a structural model with time-varying priced asset volatility to explain levels and dynamics of both credit spreads and equity volatilities. Our first contribution is to show that, in calibrations, a reasonable unlevered asset variance risk premium (AVRP) allows our model to match spread levels for medium and longer maturities without difficulty. Our second contribution is to estimate a stochastic volatility jump-diffusion (SVJ) model on firm-level data for default swap spreads and equity volatility. We identify an economically significant AVRP and find that modeling priced stochastic asset volatility strongly improves the model's ability to not only account for time variation in equity volatility, but also explain the time series of default swap spread term structures.

Recent empirical work on default swap spreads provides evidence suggestive of an important role for stochastic and priced asset volatility in credit risk modeling.[2] Although a compelling extension to a class of models that has been around for almost 40 years, stochastic volatility (SV) has not garnered much attention in the credit risk literature. We present semi–closed-form (up to a Fourier Inversion) solutions to debt and equity prices in a stochastic asset volatility framework where default is triggered by a default boundary, as in Black and Cox (1976), Longstaff and Schwartz (1995), and Collin-Dufresne and Goldstein (2001).[3] In a model with SV, different aspects of variance dynamics influence credit spreads. However, by far, the strongest effect arises from the market price of asset volatility risk.

To illustrate this point, we replicate the Huang and Huang (2012) and Chen Collin-Dufresne and Goldstein (2009) calibrations. We confirm that, in the absence of stochastic asset risk, our model replicates the credit spread puzzle. Without a risk premium on asset variance, the model faces the same problem that past studies have struggled with—it is hard to generate sufficiently high spreads. The presence of stochastic variance per se does not help sufficiently in

[2] Huang and Zhou (2008) test a broad set of structural models and show the models' inability to fit the dynamics of credit default swap prices *and* equity volatilities. In particular, they find that the models have difficulty generating sufficient time variation in equity volatility, suggesting that an extension allowing for stochastic asset volatility is desirable. Zhang, Zhou, and Zhu (2009) perform an empirical study of the effect of volatility and jumps on default swap prices. Their results also point to the importance of modeling time-varying volatility. Further evidence is provided in Wang, Zhou, and Zhou (2013), who show that not only are equity variance levels important for the price of default protection, but the associated risk premium is also a key determinant of firm-level credit spreads. Given this evidence, financial leverage would have to be the sole source of variation in stock return volatility for asset volatility to be constant, as it is assumed to be in the majority of structural credit risk models. Instead, recent empirical work by Choi and Richardsson (2016) documents time variability in unlevered asset risk.

[3] Heston (1993) provides a closed-form solution for the price of a European option with Cox, Ingersoll, and Ross's (1985) dynamics for the variance. Fouque et al. (2003, 2004, 2011) introduce perturbation techniques to address SV in a variety of option pricing settings. Related to our work, Fouque et al. (2006) use a slow variation asymptotic approximation to a free boundary problem with Gaussian variance dynamics to study defaultable securities. In the Internet Appendix, we discuss these and other methodologies. The Internet Appendix may be found in the online version of this article.

doing so. However, for reasonable parameter values governing the AVRP, our model has no difficulty matching medium- to long-term spread levels. Because short-term spreads remain hard to explain, consistent with previous work, we introduce a jump component to our asset value dynamics using a specification similar to Pan (2002) that allows for a risk premium on jump-size uncertainty.[4] To better understand this extension to our model, we conduct a calibration exercise for a representative firm as we do for the SV case. We fix Sharpe ratios and total volatility in our calibration. Note that introducing priced jumps affects the amount of variance risk the model will bear—this would not matter if jump and volatility risk were substitutes, but they influence different parts of the term structure. We find a combination that does well in explaining spreads across the term structure while allowing the model to fit both long- and short-term default probabilities.

In addition to conducting calibration exercises for representative firms, we run firm-by-firm estimations. In particular, we estimate firm-specific variance risk premia using time series of default swap term structures and option-implied volatilities for a cross section of firms with a variety of industry and credit rating characteristics. The estimations require identification of the dynamics of asset values, volatilities, and jumps, all of which are unobservable. We estimate physical and risk-adjusted firm value, volatility, and jump dynamics using a simulated maximum likelihood (ML) methodology.

To the best of our knowledge, we provide the first estimates of firm-level asset variance and jump dynamics together with their risk premia in the literature. We obtain a broad range of estimates for the AVRP that are negative and statistically significant for a large majority of firms. The implied mean equity variance risk premium (EVRP) is −16.7%, with 5% and 95% percentiles of −44% and −1.5%, respectively. The average jump intensity estimate across firms is 1%, with a cross-sectional variation of 0.1% to 4.6%, which is plausible given the variation in credit ratings in our sample. The risk-adjusted jump size is significant for the majority of firms, with an average of −81% of asset value.[5] It is important to note that, in addition to explaining credit spread levels, our model fits the level of equity volatility. Likelihood ratio tests show that our SV model strongly dominates the nested constant volatility case.

To benchmark the level of our estimates, we translate them into an EVRP that we compare with extant literature. The closest paper to our specification is Pan (2002), who produces equity index-level estimates. Given the diversification provided by an index, estimates for the index variance risk premium

[4] Since the advent of reduced-form credit models (see, e.g., Duffie and Singleton (1999)) that are based on jumps-to-default, it has been noted that diffusion-based models generate insufficient short-term spreads. Duffie and Lando (2001) find that jumps-to-default arise endogenously in a structural model when asset values cannot be observed precisely and that this helps generate significant short-term spreads. Collin-Dufresne, Goldstein, and Yang (2012) introduce jumps into a structural model for credit index spreads for precisely this reason.

[5] In the Internet Appendix, we report on the estimation of the nested SV case and find that the inclusion of jumps does not significantly change the estimates of the parameters of the volatility dynamics.

provide a conservative benchmark for what constitutes a "high" firm-level EVRP, all else equal. If idiosyncratic volatility risk is not priced, then for a given level of volatility and leverage, the EVRP for an individual firm with some idiosyncratic risk will be lower than that of the market. With this in mind, a comparison with the estimates in Pan (2002) suggests that our estimates are economically reasonable.[6]

Since we identify both physical and risk-adjusted asset value and volatility dynamics, we are able to compute physical default probabilities and Sharpe ratios. These metrics also suggest that our AVRP estimates are reasonable. Given the estimates of both physical and risk-adjusted dynamics for firm value, we decompose credit spreads into a component due to average losses, a risk premium for asset variance, and a risk premium for jump risk. The contribution of the variance risk premium to five-year credit default swap (CDS) spreads is on average 35%. Shutting down the variance risk premium significantly reduces median spreads. We also find that, in relative terms, the contribution of variance risk to spreads is larger for more highly rated firms.

The time-series fit of our model is good overall. The median bias in spreads across maturities is highest for the one-year tenor at five basis points. Most of the improvement over a traditional constant volatility model such as Black and Cox (1976) comes from the variance risk premium, while jumps help at the short end of the spread curve. Root mean squared errors are improved for all maturity segments in our SVJ specification, but remain high for the shortest maturities. Illiquidity may be a reason for the difficulty that SVJ models face in explaining short-term credit spreads.[7] To address this possibility, we reestimate the SVJ model with an adjustment for bid-ask spreads. Relative bid-ask spreads for one-year spreads are almost three times higher than those for three-year spreads. Following Pan and Singleton (2008), we allow spread observation errors to be proportional to firm-level time series of relative bid-ask spreads for each maturity. This extension yields an improvement in model fit, although not for shorter maturities.

Our baseline model takes the default boundary to be an exogenous function of leverage. We provide an extension of our SV specification to allow for endogenous default and benchmark it to the constant volatility case discussed in Leland (1994). We find that the sensitivity of the default boundary to the level of volatility depends on the parameters governing the speed of mean reversion and volatility of asset variance. Most importantly, a variance risk premium increases the value of the upside to shareholders and thus lowers the default boundary. Our variance risk premium is specified as proportional to the level

[6] Another metric to benchmark our estimates of the EVRP for the SV case is the ratio of option-implied to realized volatilities, which is an empirical measure of the variance risk premium. Han and Zhou (2011) report one-month implied and realized volatilities for a panel of 5,000 individual stocks. The ratio of their reported averages is 1.17. For our benchmark firm, the ratio of one month expected risk adjusted to physical volatility is 1.04.

[7] See, for example, Bongaerts, de Jong, and Driessen (2011) for evidence that CDS spreads may include a liquidity risk premium earned by the seller of protection.

of asset variance. Hence, our default boundary will be lower when risk premia are high. A procyclical default boundary is consistent with Chen (2010) and Bhamra, Kuehn, and Strebulaev (2010). However, the effect of variance risk on default probabilities dominates the effect of the procyclical boundary, and thus, default rates remain countercyclical, as in those papers. In addition, given that variance risk influences the decision to default, it is natural to ask how it affects firms' ex ante leverage decisions. The effect on the capital structure decision is to reduce the optimal leverage. In our stylized example, the reduction is by an economically significant amount. This finding suggests that variance risk premia should be an empirical determinant of the cross section of corporate leverage ratios.

Finally, our work relates to the recent literature on the ability of equilibrium asset pricing models to explain the credit spread puzzle. Motivated by the possibility that the equity premium and credit spread puzzles are related, CCG bring the habit preferences proposed by Campbell and Cochrane (1999) to bear on credit spreads.[8] The strong time variation in the risk premium inherent to this model helps explain credit spread levels. A parallel can be drawn to our setting, where it is the time variation in systematic asset variances, and thus asset Sharpe ratios, that allows the model to explain spreads.[9]

In Merton (1974), bondholders provide asset risk insurance to shareholders by selling a put. In our model, shareholders benefit further from (systematic) volatility insurance written by bondholders. This is reflected in lower expected instantaneous returns and Sharpe ratios for equity of highly levered firms. This observation is consistent with recent evidence of low equity returns for firms near distress (Campbell, Hilscher, and Szilagi (2008)).[10] On the other hand, since bondholders are the insurance writers, they require higher risk premia to compensate. Accordingly, in our model, debt Sharpe ratios are higher than

[8] Consistent with this insight, Chen (2010) and Bhamra, Kuehn, and Strebulaev (2010) show that the long-run risk framework of Bansal and Yaron (2004) with regime-switching expected growth rates and volatility combined with a dynamic trade-off model can generate sufficient co-movement in risk premia, default rates, and default losses to explain the high credit spreads and low leverage ratios of investment-grade firms.

[9] CCG further find that a combination of habit preferences with countercyclical asset value default boundaries can resolve the credit spread puzzle. Our model is able to generate sufficient spread levels with either a fixed or endogenously procyclical boundary, a feature shared by the models of Chen (2010) and Bhamra, Kuehn, and Strebulaev (2010).

[10] McQuade (2018), who, like us, calibrates a structural model of the firm with priced stochastic asset volatility to study credit spreads, also finds that his model can predict lower equity returns near distress. His model has predictions relevant for the momentum anomaly and the value premium. He finds that allowing for endogenous default helps his model fit credit spreads in the cross section of credit ratings. He does not, however, carry out firm-by-firm estimation of his model, nor does he consider optimal capital structure. Garlappi, Shu, and Yan (2008) show that in a one-factor model in which shareholders and bondholders bargain about the proceeds in distress, the equity beta decreases near distress, providing a possible solution to the puzzle. With variance risk, the regular equity beta can be monotonically increasing in leverage but is offset by an increased loading on the negatively priced variance risk.

their equity counterparts.[11] This points to a possible link between the distress puzzle and the credit spread puzzle.

This paper is organized as follows. Section I describes the SV model and explains how we derive it in closed form. Section II covers calibration of the model. Section IV describes our SVJ model, discusses comparative statics, and presents calibration results. Section V reports the estimation results for the SVJ model. We provide further extensions to our model in Section VI and conclude in Section VII.

I. Model

To study the role of variance risk in pricing default risk, we extend the Leland (1994) model to incorporate SV. This involves variance dynamics like those in Heston (1993), interpreted as asset rather than stock return variance. Asset value, X, and variance, V, dynamics are described as follows under the physical probability measure P:

$$\frac{dX_t}{X_t} = \left(\mu_X^P - q\right)dt + \sqrt{V_t}\left(\rho dW_{2,t} + \sqrt{1-\rho^2}dW_{1,t}\right),$$
$$dV_t = \kappa(\theta - V_t)dt + \sigma\sqrt{V_t}dW_{2,t},$$

where μ_X^P is the expected return on the firm's assets, q is the proportional payout rate, ρ is the instantaneous correlation between the Brownian motion $W = \rho W_2 + \sqrt{1-\rho^2}W_1$ that drives asset value uncertainty and W_2 that drives uncertainty with respect to asset variance, κ is the speed of mean reversion, θ is the long-run mean variance, and σ is the volatility parameter for asset variance. This specification for unlevered asset dynamics involves both diffusive risk and uncertainty about future asset variance. To reflect this, we specify the mean of the unlevered asset return μ_X^P as a function of the risk premia for the two sources of risk. More specifically, we let the expected unlevered asset price appreciate with the risk-free rate r and two risk premium components: $\mu_X^P = r + (\sqrt{1-\rho^2}\lambda_D + \rho\lambda_V)V_t$. This expected return compensates for diffusive risk via λ_D and variance risk via λ_V. The risk-adjusted dynamics for X and V are

$$\frac{dX_t}{X_t} = (r-q)dt + \sqrt{V_t}\left(\rho dW_{2,t}^Q + \sqrt{1-\rho^2}dW_{1,t}^Q\right),$$
$$dV_t = \kappa^*\left(\theta^* - V_t\right)dt + \sigma\sqrt{V_t}dW_{2,t}^Q,$$

[11] In a model in which uncertainty is normally distributed, the Sharpe ratio is the same for debt and equity, regardless of leverage. In our model, this is not the case and Sharpe ratios differ depending on the mix of loadings on the asset level and variance risks.

where $\kappa^* = (\kappa + \sigma\lambda_V)$ and $\theta^* = \frac{\kappa}{\kappa^*}\theta$. Finally, a quantity that we discuss often below is the asset return Sharpe ratio, which is given by

$$SR_t = \frac{\mu_X^P - r}{\sqrt{V_t}} = \left(\sqrt{1-\rho^2}\lambda_D + \rho\lambda_V\right)\sqrt{V_t}. \tag{1}$$

In what follows, we describe the solution for the firm's equity value and equity volatility. As in Leland (1994), we assume that the firm issues consol bonds. The equity value can be written as the difference between the levered firm value (F) and the debt value (D), that is, $E(X) = F(X) - D(X)$. The levered firm's value is given by

$$F(X_t) = X_t + \frac{\zeta c}{r}(1 - p_D(X_t, V_t)) - \alpha X_D p_D(X_t, V_t), \tag{2}$$

where X, ζ, c, α, X_D, and p_D, respectively, denote the initial unlevered asset value, the tax rate, the coupon rate, the liquidation cost, the default boundary, and the present value of \$1 at default. At this stage, we take the default boundary as an exogenously given constant. In Section VI, we provide an extension in which the firm's default policy is endogenous. In equation (2), the first term is the unlevered asset value, the second term is the tax benefit, and the third term is the bankruptcy cost. The debt value is the sum of the present value of the coupon payments before default and the recovered value of the firm at default, and is given by

$$D(X_t) = \frac{c}{r} + \left[(1-\alpha)X_D - \frac{c}{r}\right] p_D(X_t, V_t).$$

The equity value is therefore given by

$$E(X_t) = X_t - \frac{(1-\zeta)c}{r} + \left[(1-\zeta)\frac{c}{r} - X_D\right] p_D(X_t, V_t). \tag{3}$$

Applying Itô's lemma, we obtain the stochastic process for the equity value as follows:

$$\frac{dE_t}{E_t} = \mu_{E,t}dt + \frac{X_t}{E_t}\frac{\partial E_t}{\partial X_t}\sqrt{V_t}\left(\rho dW_{2,t} + \sqrt{1-\rho^2}dW_{1,t}\right) + \sigma\frac{1}{E_t}\frac{\partial E_t}{\partial V_t}\sqrt{V_t}dW_{2,t},$$

where $\mu_{E,t}$ is the instantaneous equity return. The instantaneous equity volatility, $\sigma_{E,t}$, can be written as

$$\sigma_{E,t} = \sqrt{\left[\left(\frac{X_t}{E_t}\frac{\partial E_t}{\partial X_t}\right)^2 + \left(\frac{\sigma}{E_t}\frac{\partial E_t}{\partial V_t}\right)^2 + 2\rho\sigma\frac{X_t}{E_t^2}\frac{\partial E_t}{\partial X_t}\frac{\partial E_t}{\partial V_t}\right]V_t}. \tag{4}$$

Under the risk-adjusted measure Q, $p_D = E_t^Q[e^{-r(\tau_D - t)}]$, where $\tau_D = \inf\{s \geqslant t \mid X_s = X_D\}$ denotes the default time. To solve for p_D, we need to compute the probability density function of the stopping time τ_D under the risk-adjusted

probability measure Q. To do so, we rely on Fortet's lemma, as discussed in the Internet Appendix.[12]

II. Calibrating the SV Model

Prior to carrying out a full estimation of the model, we consider calibrated examples to illustrate its key implications. The credit spread puzzle is typically seen as an inability to generate a sufficiently large risk premium component for bond spreads, in a model calibrated to default losses, leverage, and some measure of the equity risk premium.[13] CCG show that a defaultable bond price can be expressed as a function of the physical default probability, the expected loss, and the Sharpe ratio. This suggests that in a calibration where the first two are fixed, the Sharpe ratio will determine the attainable spread levels. In fact, they show that if most asset risk is systematic, it is not hard to fit historical spreads. As a result, they advocate calibrating models to Sharpe ratios as well to account for the presence of both idiosyncratic and systematic volatility.

Like CCG, we consider an initial state of the world based on historical asset volatility, Sharpe ratio, and default probability for a representative Baa firm.[14] We assume an annualized asset volatility of 29%, a physical default probability of 4.9% at the 10-year horizon, and a historical Sharpe ratio of 22%.[15] The asset value drift is determined by the asset variance and the Sharpe ratio defined in equation (1). As we will show, this constraint, together with the matching of historical default rates, significantly limits the range of economically viable values for the amounts of diffusive and variance risk premia.

We set the correlation between asset value and variance shocks to -0.15, the volatility of asset variance (σ) to 30%, the long-run mean variance (θ) equal to the base-case initial variance (0.29^2), and the speed of mean reversion (κ) to 4. There are no empirical estimates to benchmark the firm-by-firm asset volatility parameters. In the absence of any clear comparable, we justify the levels for the speed of mean reversion (κ), the correlation between asset value and variance shocks (ρ), and the volatility of variance (σ) by the levels we obtain when we estimate the model in Section V. The speed of mean reversion is somewhat lower than the level found in Aït-Sahalia and Kimmel (2007) for equity index option data. Estimates for the correlation parameter available at the equity index level suggest that ρ lies between -0.4 and -0.5 (see, e.g., Eraker, Johannes, and Polson (2003)). Asset value correlation may be lower

[12] See Fortet (1943).

[13] See Huang and Huang (2012).

[14] Historical leverage ratios, payout rates, and recovery rates are also used as inputs.

[15] Asset volatility of 29% is based on the median estimate in Feldhütter and Schaefer (2018) for that rating category. We use the Sharpe ratio reported in CCG for the period 1974 to 1998. While in Merton (1974), Sharpe ratios of equity and debt are identical; later we show that in our setting they may differ. Equity Sharpe ratios are lower than asset Sharpe ratios. For the purposes of this particular calibration the difference is not economically significant and by constraining the somewhat higher asset Sharpe ratio, we are being conservative in terms of generating spreads. The source for the default probabilities is Moody's report for the period 1970 to 2010.

for several reasons. First, part of the negative relation for equity returns may be related to financial leverage. Second, there is evidence that the asymmetric volatility effect is lower for idiosyncratic volatility shocks than for their systematic counterparts (see Dennis, Mayhew, and Stivers (2006)). Estimates in the equity index option literature for the volatility of variance ranges from around 0.3 (e.g., Bates (2006) and Pan (2002)) to 0.5 (e.g., Aït-Sahalia and Kimmel (2007)).

To adjust for illiquidity, we follow CCG and consider the spread to Aaa as a proxy for the default component of a firm's bond yield. This way of accounting for nondefault factors implicitly assumes that it is mostly the default component of spreads that varies with the credit rating. In the Internet Appendix, we provide an alternative adjustment for illiquidity based on bid-ask spreads. The historical average for the 10-year Baa-Aaa spread is 131 basis points.

Table I reports on the calibration for our representative Baa firm. Consider first line 4, where the initial asset volatility is set to a historical level of 29%. To illustrate the influence of variance risk premia on spreads, two polar scenarios are considered: one with the risk premium determined entirely by the variance risk premium ($\lambda_D \equiv 0$) and the other entirely by a diffusive asset risk premium ($\lambda_V \equiv 0$). We also consider a combination of variance and diffusive risk that matches the historical Baa-Aaa spread for an initial asset volatility of 29%. Given the level of drift and asset volatility, we match 10-year default probabilities by selecting the level of the default boundary.[16]

The credit spread puzzle has been presented in past studies as the inability to generate, for example, Baa-Aaa spreads near their historical levels. For instance, in Huang and Huang (2012), model-implied 10-year Baa spread levels rarely exceed 60 basis points across their selection of models.[17] At the benchmark asset volatility (line 4), we obtain a spread of 84 basis points without the variance risk premium. With priced variance risk, our model-implied spread reaches 183 basis points at the benchmark level of asset volatility. This is more than twice the spread without variance risk, it is higher than the historical 10-year Baa-Aaa spread and, it is slightly lower than the 10-year Baa to Treasury spread of 194 basis points reported in Huang and Huang (2012). Note, however, that these results are based on using the maximum amount of variance risk premium consistent with the historical Sharpe ratio. More weight given to the diffusive risk premium would reduce the predicted spread. We carry out a similar exercise for four-year spreads and obtain similar results.

In summary, Table I allows us to make two important observations. First, without a variance risk premium, we face the same problem that past studies have struggled with—it is hard to generate sufficiently high spreads. The

[16] To match 10-year probabilities, the default boundary equals 43.5% of the total liabilities. This value is in the range reported in Davydenko (2012), although lower than the median of 61% of debt face value.

[17] For their base-case formulation using the Longstaff and Schwartz (1995) model under a constant risk free rate, they report 56 basis points. We will refer to this base-case model whenever we cite their findings.

Table I

Variance Risk Premia and Credit Spreads for a Representative Baa Firm

This table considers three different risk premium compositions: The two polar cases where all risk compensation arises from the variance risk (*Variance risk only* column) or diffusive risk (*Diffusive risk only* column), and a mix between the two that matches the historical Baa-Aaa 10-year spread. The model is calibrated to match a Sharpe ratio of 22% and cumulative default probability of 4.9% at the 10-year horizon, at the initial volatility level of 29%. Hence, the fourth row represents the initial state. We consider how default probabilities, spreads, and Sharpe ratios vary with the level of asset risk. The risk-free rate is 5%, the recovery in default is 51%, and the payout rate is 5%. The volatility dynamics under P are described by the initial variance, a volatility of volatility (σ) of 30%, a correlation (ρ) between the asset level and variance shocks of -0.15, and a speed of mean reversion (κ) of 4. The default boundary is chosen to match the default rate. When $\lambda_V \equiv 0$, we have $\lambda_D = 0.77$, and when $\lambda_D \equiv 0$, we have $\lambda_V = -5.06$. The mix that matches the historical spread of 131 basis points is $\lambda_D = 0.3$, $\lambda_V = -3.08$. PD and QD denote the default probability under the physical and risk-adjusted measures, respectively.

				Variance Risk Only				Diffusive Risk Only				λ_D^*, λ_V^*			
Initial Vol ($\sqrt{V_0}$)	Sharpe Ratio	PD 1Y	PD 10Y	QD 1Y	QD 10Y	Spread 1Y	Spread 10Y	QD 1Y	QD 10Y	Spread 1Y	Spread 10Y	QD 1Y	QD 10Y	Spread 1Y	Spread 10Y
0.20	0.15	0.000	0.007	0.000	0.091	0	49	0	0.025	0	13	0.000	0.053	–	28
0.23	0.17	0.000	0.018	0.000	0.152	0	84	0	0.056	0	30	0.000	0.103	–	55
0.26	0.20	0.000	0.031	0.000	0.231	1	133	0	0.101	0	54	0.000	0.163	–	90
0.29	0.22	0.000	0.049	0.000	0.304	1	183	0	0.152	0	84	0.000	0.229	1	131
0.32	0.24	0.000	0.068	0.000	0.387	5	244	0	0.212	1	121	0.000	0.297	3	178
0.35	0.27	0.000	0.089	0.000	0.454	7	305	0	0.270	2	162	0.000	0.371	7	232
0.38	0.29	0.000	0.107	0.000	0.520	23	375	0	0.339	3	208	0.000	0.434	13	287

presence of SV per se does not help in doing so.[18] Second, in the presence of maximal variance risk, predicted 10-year spreads are much higher, and in fact overshoot required levels. More importantly, there exists a mix between diffusive and variance risk ($\lambda_D = 0.3$ and $\lambda_V = -3.08$) that is able to match the historical Baa-Aaa spread of 131 basis points.

Having calibrated the model to the 29% volatility case, we consider alternative scenarios for this metric. Intuitively, we ask what would happen if at some point in the near future we faced a new volatility level ranging from 20% to 38%, while all other parameters and initial asset value were held constant. A key feature of our model is that the asset value drift and hence Sharpe ratio depends on the level of the variance. When variance is high, the compensation per unit asset risk is high, and vice versa. This is evident from the second column in Table I: Across the range of asset volatilities between 20% and 38%, the Sharpe ratio varies between 15% and 29%. This is reminiscent of the findings in CCG, where the countercyclical compensation for risk helps generate higher spreads.[19]

In our specification of the asset drift, the Sharpe ratio is state dependent regardless of the composition of the asset risk premium $\mu_X^P - r$. However, note that the wedge between risk-adjusted and physical default probabilities widens much more in high volatility states in the presence of a variance risk premium. For example, while the asset Sharpe ratio increases from 22% to 29% for both risk premium specifications when asset volatility rises to 38%, the ratio between Q and P 10-year default rates is about 5 with variance risk only, compared to about 3 without. This translates into a spread increase of 192 basis points compared to 124. At λ_D^*, λ_V^*, the ratio is about 4 and the spread increases by 156 basis points.[20] It therefore appears that while constrained to generate reasonable Sharpe ratios and overall volatilities, our variance risk specification also allows us to generate reasonable medium- to long-term credit spread levels.

We now turn to a typical Aa firm in Table II. This exercise is justified by the observation (e.g., Huang and Huang (2012)) that the credit spread puzzle becomes more severe as credit quality increases. Table II reports on an Aa firm with 21% asset volatility and 21% leverage.[21] With diffusive risk only, we find a spread of 18 basis points, whereas with variance risk only, we obtain 99 basis points. This compares to Huang and Huang (2012), who find a 14 basis point 10-year spread implied in their base-case calibration using a constant-interest-rate version of Longstaff and Schwartz (1995). They report a historical

[18] As pointed out by CCG, once a model is calibrated to expected losses, spreads can only be influenced by the covariance between the pricing kernel and future cash flows.

[19] See Lustig and Verdelhan (2012) for empirical evidence of the countercyclicality of equity Sharpe ratios. We provide a more detailed discussion of equity and asset return Sharpe ratios in the Internet Appendix.

[20] These ratios are comparable to what has been reported in Berndt et al. (2018), and Doshi et al. (2013).

[21] The asset volatility value is based on Feldhütter and Schaefer (2018) and the leverage value from Huang and Huang (2012).

Table II

Variance Risk Premia and Credit Spreads for a Representative Aa Firm

This table considers three different risk premium compositions: the two polar cases where all risk compensation arises from either variance risk or diffusive risk, respectively, and a mix between the two (denoted $(\lambda_D^*, \lambda_V^*)$) that matches the historical Aa-Aaa 10-year spread. The model is calibrated to match a Sharpe ratio of 22% and cumulative default probability of 0.6% at the 10-year horizon, at the initial volatility level of 21%. Hence, the fourth row represents the initial state. We consider how default probabilities, spreads, and Sharpe ratios vary with the level of asset risk. The risk-free rate is 5%, the recovery in default is 51%, and the payout rate is 5%. The volatility dynamics under P are described by the initial variance, a volatility of volatility (σ) of 30%, a correlation (ρ) between the asset level and variance shocks of -0.15, and a speed of mean reversion (κ) of 4. The default boundary is chosen to match the default rate. When $\lambda_V \equiv 0$, we have $\lambda_D = 1.06$, and when $\lambda_D \equiv 0$, we have $\lambda_V = -6.98$. The mix that matches the historical spread of 28 basis points is $\lambda_D = 0.77$, $\lambda_V = -1.93$. PD and QD denote the default probability under the physical and risk-adjusted measures, respectively.

				Variance Risk Only				Diffusive Risk Only				λ_D^*, λ_V^*			
Initial Vol ($\sqrt{V_0}$)	Sharpe Ratio	PD 1Y	PD 10Y	QD 1Y	QD 10Y	Spread 1Y	Spread 10Y	QD 1Y	QD 10Y	Spread 1Y	Spread 10Y	QD 1Y	QD 10Y	Spread 1Y	Spread 10Y
0.12	0.13	0.000	0.000	0.000	0.014	0	7	0.000	0.000	0	0	0.000	0.001	0	0
0.15	0.16	0.000	0.000	0.000	0.046	0	25	0.000	0.003	0	1	0.000	0.006	0	3
0.18	0.19	0.000	0.002	0.000	0.102	0	56	0.000	0.011	0	6	0.000	0.020	0	11
0.21	0.22	0.000	0.006	0.000	0.176	0	99	0.000	0.035	0	18	0.000	0.053	0	28
0.24	0.25	0.000	0.013	0.000	0.262	1	154	0.000	0.072	0	38	0.000	0.100	0	54
0.27	0.28	0.000	0.023	0.000	0.356	2	219	0.000	0.117	0	64	0.000	0.156	0	86
0.30	0.31	0.000	0.032	0.000	0.432	6	286	0.000	0.173	0	97	0.000	0.220	1	126

spread to Treasury of 91 basis points and a spread to Aaa of 28 basis points. For an Aa firm, it is also feasible to determine a combination λ_D^*, λ_V^* to match either historical spread metric. Interestingly, ratios of risk-adjusted to physical 10-year default probabilities are higher (for any combination of diffusive and variance risk) for an Aa firm than a Baa firm.[22]

It has often been observed that, due to the predictability of default, structural models are unable to generate nontrivial short-term spreads. In addition, we are aware of little work that discusses the ability of structural models to explain short-term physical default probabilities. An exception is Leland (2006), who documents a tendency to underestimate not only spreads but also default rates at shorter horizons. Table I reports on spreads and default probabilities at the one-year horizon and confirms the current model specification's inability to generate sufficient spreads and default rates one year out. This is an interesting variation of the credit spread puzzle, which most often refers to high empirical spreads with low empirical default rates. At the short end, the challenge is to generate sufficiently high spreads and default rates, whereas at longer horizons, the issue is to fit spreads without exaggerating default rates. We return to this tension when introducing jumps into our framework below.

In summary, introducing a variance risk premium significantly increases medium- to long-term spreads, for both Baa and Aa firms, while matching historical default probabilities, Sharpe ratios, and volatilities. A key feature of our model is that the Sharpe ratio is state dependent, as is the difference between risk-adjusted and physical default probabilities. The Sharpe ratio is high when volatility is high and increases spreads significantly more in high volatility states than in the absence of a variance risk premium. The current model specification does not improve, however, on past models' performance at shorter horizons.

One difficulty we face in the analysis above is that we do not have empirical estimates for the AVRP parameter λ_V. In the options literature, the EVRP parameter has been estimated for stock indices. For example, Pan (2002) finds a value equal to 7.6 using data on short-dated S&P 500 index options. Unfortunately, this value is not directly comparable to our parameter λ_V for two reasons. First, our risk premium applies to individual firms as opposed to indices and, second, it reflects the risk of unlevered assets rather than levered equity. To conduct a more appropriate comparison, we derive the relationship between asset and equity variance risk premia.

More specifically, we measure the variance risk premium using the change in the drift of the variance dynamics when moving from the physical to the risk-adjusted probability measure. For unlevered variance dynamics, this change in drift is given by $\lambda_V \sigma\ V_t dt$. This translates into the following expression for the EVRP:

$$EVRP(X_t, V_t) = \left(\frac{\partial V_{E,t}}{\partial V_t} \sigma \lambda_V + X_t \frac{\partial V_{E,t}}{\partial X_t} \left(\rho \lambda_V + \sqrt{1-\rho^2} \lambda_D \right) \right) V_t. \tag{5}$$

[22] See Berndt et al. (2018).

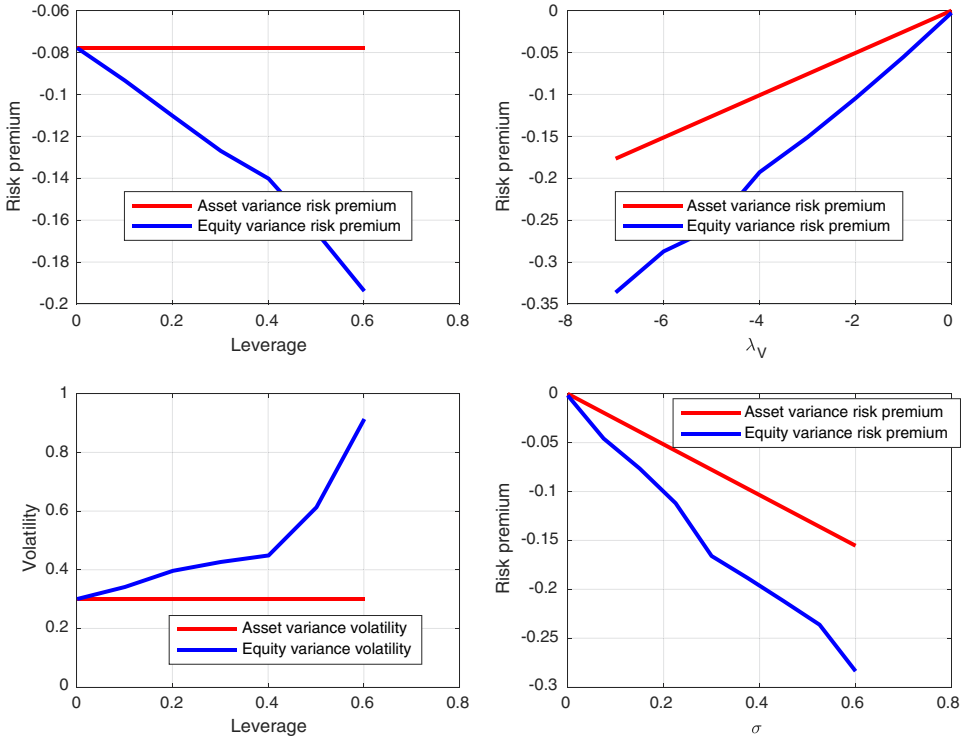

Figure 1. Levered and unlevered variance risk premia and volatilities. This figure illustrates the interplay between leverage, volatility of assets, and the equity variance risk premia (EVRP), computed using equation (5). Panel A plots asset and equity variance risk premia versus leverage. Panel B plots the two types of variance risk premia against the variance risk premium parameter λ_V. Panel C plots the volatility of asset and equity variance dynamics against leverage. Finally, Panel D plots variance risk premia against the volatility of asset variance. (Color figure can be viewed at wileyonlinelibrary.com)

Note that the EVRP is a function of both the diffusive and the variance risk parameters, λ_D and λ_V.[23]

The EVRP is a function of all the variables and parameters of the model. Figure 1 summarizes selected comparative statics of equity versus asset variance premia across leverage, variance volatility, and the AVRP parameter λ_V. The top left panel illustrates how a firm's EVRP is amplified by leverage (without leverage it coincides with the AVRP). With an AVRP of -7.5% per annum, its levered counterpart is almost doubled at just less than -15% when leverage reaches 40% and almost trebled once it reaches 60%. As we discuss below, this will have a tendency to reduce equity returns, all else equal.

While, as expected, both risk premium metrics coincide at zero, as λ_V becomes more negative (upper right panel), both risk premium metrics become

[23] In the Internet Appendix, we provide details on the derivation of the EVRP in the more general SVJ specification.

Table III

The Equity versus Asset Variance Risk Premium

The model is calibrated to match a Sharpe ratio of 22% and cumulative default probability of 4.9% at the 10-year horizon, and at the initial volatility level of 29%. The risk-free rate is 5%, the recovery in default is 51%, and the payout rate is 5%. The volatility dynamics under P are described by the initial variance, a volatility of volatility (σ) of 30%, a correlation (ρ) between the asset level and variance shocks of -0.15, and a speed of mean reversion (κ) of 4. The default boundary is chosen to match the default rate at 10 years. We set the risk premium parameters to the levels that match the historical 10-year Baa-Aaa spread of 131 basis points, implying $\lambda_D = 0.30$, $\lambda_V = -3.08$. All columns except spreads are in units of percentage points; spreads are in basis points. Panel A provides asset and variance risk premia for three asset volatility scenarios, corresponding to the calibration in Table I for a representative Baa firm. Panel B reports the levered index variance risk premium implied by the parameter estimates in Pan (2002).

Panel A: Individual firm-level levered and unlevered variance risk premium

Scenario	Initial Asset Volatility	Sharpe Ratio	Leverage	Asset VRP	Equity VRP
Low vol.	0.20	0.15	0.45	−0.0370	−0.0552
Baseline	0.29	0.22	0.45	−0.0780	−0.1091
High vol.	0.38	0.29	0.45	−0.1334	−0.1856

Panel B: S&P 500 index level variance risk premium

Initial index volatility	Index VRP
15	−0.17
20	−0.304

more negative as well. The EVRP, however, does so much faster, reflecting its embedded leverage. The lower left panel translates the volatility of unlevered variance (σ) into the volatility of equity variance. With an asset variance volatility of 30% per annum, the equity variance volatility ranges from 30% at zero leverage to more than 80% at 60% leverage. As the volatility of variance increases, a similar pattern is observed. Both metrics begin at zero and become more negative as the volatility of asset variance increases. Overall, leverage has an economically significant effect on the translation of asset to equity variance risk premia.

Table III reports on the levels of asset and equity variance risk premia (implied from equation (5)) for the Baa firm represented in Table I. For the baseline case, where asset volatility is 29% and the Sharpe ratio is 22%, the AVRP equals about 8%, while its levered counterpart reaches just less than 11%. Given the diversification provided by an index, estimates for the index variance risk premium provide a conservative benchmark for what constitutes a "high" individual EVRP. If idiosyncratic volatility risk is not priced, then for a given level of volatility and leverage, the EVRP of an individual firm with some idiosyncratic risk will be lower than that of the market.

To better understand this claim, note that our model attaches a variance risk premium to the total asset volatility. We think of this as a reduced-form

way of letting the systematic part of a firm's asset variance be priced, while acknowledging that, as a result, the numbers that our calibration implies for λ_V are going to be lower than those implied by a model that explicitly distinguishes between unpriced idiosyncratic risk and priced systematic variance risk.

For comparison with our implied levels of individual equity variance risk premia, in Panel B, we provide the levered index variance risk premium implied by the parameter estimates in Pan (2002). The index variance risk premium ranges between −17% and −30% for index volatilities. Hence, our baseline implied variance risk premium does not seem excessive.[24] In fact, it is possible that idiosyncratic variance risk does carry a premium, in which case our λ_V would represent an average of the risk premium carried by the idiosyncratic and systematic variances. We return to the distinction between systematic and idiosyncratic variance risk in Section VI, where we discuss a model in which both risks are modeled separately.

Another measure of risk that we discuss in some detail is the Sharpe ratio. In a one-factor model like Merton (1974), the Sharpe ratios of all corporate securities and liabilities coincide, and any difference in leverage across instruments influences returns and variances in the same way. However, in our two-factor setting, this is no longer the case. Equity and debt load differently on variance risk. To see this, consider the equity Sharpe ratio in our setting,

$$\frac{\mu_E - r}{\sigma_E} = \frac{\frac{\partial E}{\partial X}(\mu_X - r)X + \frac{\partial E}{\partial V}V\sigma\lambda_V}{\sigma_E E}. \tag{6}$$

In Merton (1974), creditors are short a put enjoyed by shareholders. In our model, the value of this put is driven by both asset risk and variance risk. One can think of this as shareholders benefiting from systematic variance insurance sold by bondholders. The loading on the (negative) variance risk premium in the equity Sharpe ratio, equation (6), is positive, which tends to reduce the return required by shareholders all else equal. While leverage amplifies asset variance and the required return on assets, it increases the loading on variance risk, dampening the effect of leverage.[25] However, there is no such offsetting effect on the equity variance in the denominator, which leads to a lower Sharpe ratio for higher levels of leverage. The opposite is true for debt. In the analogue of equation (6) for debt, the variance risk premium has the opposite effect, increased leverage amplifies returns faster than debt variance, and hence equity Sharpe ratios are lower than their debt counterparts. This is illustrated in Figure 2. The lower equity Sharpe ratio is reminiscent of the distress puzzle, whereby returns on stocks for firms near financial distress

[24] We have also computed the ratio of expected risk-adjusted to physical volatility at the one-month horizon for the benchmark firm in Table I and find it to be 1.04. This can be compared with the ratio implied by Han and Zhou (2011) for a large sample of 5,000 individual stocks, which is higher at 1.17.

[25] For a sufficiently high AVRP, this leads to a nonmonotonic relationship between leverage and equity returns.

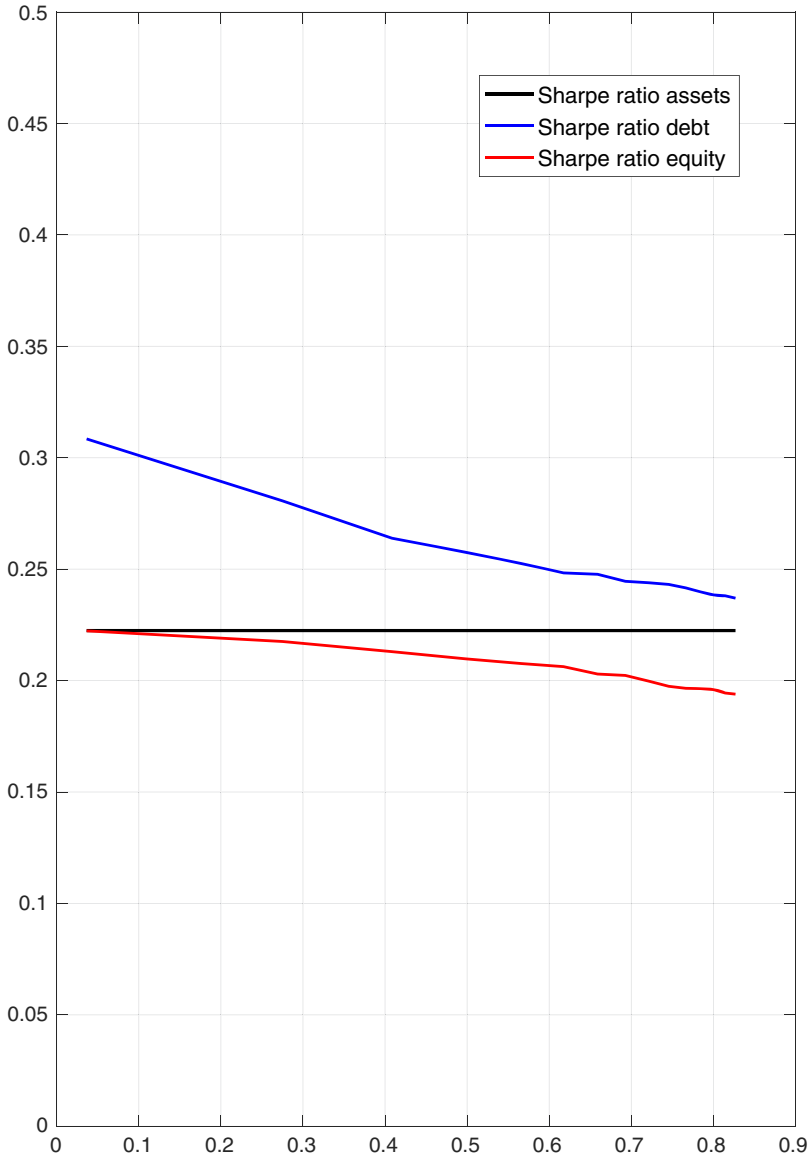

Figure 2. Sharpe ratios in the presence of variance risk. This figure plots the variation in Sharpe ratios of asset, equity, and debt as a function of a firm's leverage ratio. The parameters are those of the base case depicted in Table I. (Color figure can be viewed at wileyonlinelibrary.com)

appear low, given the increased risk induced by leverage. In our setting, leverage has two offsetting effects and may not increase returns.

On the other hand, the insurance sold by bondholders, which carries a risk premium, makes them require higher returns and spreads. This is consistent

with the credit spread puzzle, where a model without variance risk would tend to underestimate spreads, and suggests that these two puzzles could be related.

We now consider an alternative specification of the AVRP before discussing our model with jumps that we take to the data.

III. An Alternative Risk Premium Specification

In our benchmark model, we specify the variance risk premium as completely affine in the instantaneous variance (see Dai and Singleton (2000)). In this section, we consider an alternative, more flexible, setup where the risk premium evolves according to the extended affine specification proposed by Cheridito, Filipović, and Kimmel (2007).[26] The Internet Appendix provides a more detailed description of this setup. The impact of the change from the physical to the risk-adjusted measure is that the Brownian motion driving the variance dynamics shifts as follows:

$$dW_{2,t}^{Q} = dW_{2,t}^{P} + \left(\lambda_1\sqrt{V_t} + \lambda_2\frac{1}{\sqrt{V_t}}\right)dt. \tag{7}$$

This implies an instantaneous asset Sharpe ratio of

$$SR_t = \frac{\sqrt{1-\rho^2}\lambda_D V_t + \rho(\lambda_1 V_t + \lambda_2)}{\sqrt{V_t}}. \tag{8}$$

Different combinations of λ_1 and λ_2 can be consistent with a given initial Sharpe ratio. Table IV illustrates how varying these parameters influences risk-adjusted default probabilities and credit spreads over a 10-year horizon. The middle row, where $\lambda_2 = 0$, corresponds to the baseline scenario in Table I. A larger and more negative λ_2 requires a less negative or even positive λ_1 to maintain the Sharpe ratio at 22%. This makes the Sharpe ratio less dependent on future values for asset variance and leads to somewhat lower credit spreads, all else equal. Conversely, as λ_2 grows less negative and even positive, the required λ_1 becomes larger and more negative, making the Sharpe ratio more sensitive to future variance, thus increasing credit spreads. Letting λ_2 vary from -0.3 to 0.3 generates a range of spreads of more than 40 basis points. However, when the variance reaches low levels, a positive λ_2 will tend to generate negative Sharpe ratios. If λ_2 were negative, it would generate Sharpe ratios that are nonmonotonic in the asset variance.

[26] In the literature on affine term structure models and intensity-based credit risk models, a number of different specifications have been proposed and analyzed. For example, Duffee (2002) proposes an essentially affine risk premium specification where the price of risk and the variance of the state variable are less tightly connected than in completely affine models, while Cheridito, Filipović, and Kimmel (2007) propose an extended affine representation with further embedded flexibility. In our specific two-factor setup, the essentially affine specification reduces to the completely affine model, whereas the extended affine setting does not and provides more flexibility than our baseline setup.

Table IV

An Extended Affine Variance Risk Premium Specification and Its Effect on Credit Spreads

This table reports on the effect of varying the parameters of the extended affine risk premium specification in such a way that the Sharpe ratio is kept constant. The main parameters are those in the baseline scenario in Table I: The model is calibrated to match a Sharpe ratio of 22% and cumulative default probability of 4.9% at the 10-year horizon, at the initial volatility level of 29%. The risk-free rate is 5%, the recovery in default is 51%, and the payout rate is 5%. The volatility dynamics under P are described by the initial variance, a volatility of volatility (σ) of 30%, a correlation (ρ) between the asset level and variance shocks of −0.15, and a speed of mean reversion (κ) of 4.

QD 10Y	Spread 1Y	Spread 10Y	λ_1	λ_2	θ^*	κ^*	$\theta^*\kappa^*$
21	1	120	0.49	−0.3	0.103	4.15	0.43
22	1	125	−0.70	−0.2	0.105	3.79	0.40
22	1	129	−1.90	−0.1	0.107	3.43	0.37
23	1	131	−3.08	0	0.109	3.08	0.34
25	2	143	−4.27	0.1	0.113	2.72	0.31
26	2	153	−5.46	0.2	0.117	2.36	0.28
28	2	164	−6.65	0.3	0.123	2.01	0.25

The last three columns report on the risk-adjusted dynamics of the asset variance. As we move down the rows, and λ_1 becomes more negative, the long-run mean variance (θ^*) increases together with lower levels of mean reversion (κ^*). In the completely affine version of our model, the product of the mean reversion speed and long-run mean of the variance remains the same under both probability measures; this constraint is relaxed in the extended affine setting. Decoupling the speed of mean reversion and the long-run mean variance allows more pronounced time variation in the risk premium than under a completely affine variance risk premium, causing the model to generate higher credit spreads.[27] However, although this risk premium specification would render the model more flexible and allow higher spreads given a total variance and Sharpe ratio, we employ our original specification as the model we ultimately estimate. We find that it is very difficult to empirically identify the additional parameter, especially when our model is extended to include jumps, which we discuss next.

IV. The SVJ Model

Given the difficulties we document above for our model to fit short-term spreads, it is natural to ask (i) whether jumps and jump risk premia can help the model do so and (ii) whether the economic importance of the variance risk premium remains in such a setting. To address these questions, we next consider a jump diffusion SVJ model for the unlevered asset value.

[27] Cheridito, Filipović, and Kimmel (2007) emphasize that, to avoid arbitrage in the model, the Feller conditions are required to hold under both probability measures. In the scenario illustrated here, these conditions are not binding.

We assume the following dynamics for the unlevered asset price under P:

$$\frac{dX_t}{X_t} = \left(\mu_X^P - q - \lambda\bar{\eta}^P\right)dt + \sqrt{V_t}\left(\rho dW_2 + \sqrt{1-\rho^2}dW_1\right) + dJ_t^P$$

$$dV_t = \kappa(\theta - V_t)dt + \sigma\sqrt{V_t}dW_2,$$

where J_t^P is a jump process with constant intensity λ and a random jump size equal to η^P. Conditional on a jump, the asset value X_t jumps to $X_t exp(u^P)$, so that $\eta^P = exp(u^P) - 1$. We assume that u^P is normally distributed: $u^P \sim N(\bar{u}^P, \gamma^2)$. Following Pan (2002), we allow the risk-neutral relative jump size η^Q to be different from its physical counterpart, which allows for a premium on jump-size uncertainty. For simplicity, we abstract from a risk premium on uncertainty related to the jump intensity. The SV dynamics are the same as in the previous section.[28]

As in Section I, we specify the mean of the unlevered asset returns as a function of the risk premia for the different sources of uncertainty. The drift of the unlevered asset value now depends on the interest rate r, the payout rate q, and the three risk-premium components λ_D, λ_V, and $\lambda(\bar{\eta}^P - \bar{\eta}^Q)$, which are associated with compensation for the two Brownian risks and the jump risks, respectively.[29] Let $\mu_X^P = r + (\sqrt{1-\rho^2}\lambda_D + \rho\lambda_V)V_t + \lambda(\bar{\eta}^P - \bar{\eta}^Q)$.

In this specification with jumps, we are unable to derive a closed-form solution for the term structure of risk-adjusted default rates as we were for the SV case. We resort to Monte Carlo simulation, which we couple with function approximation to make the estimation tractable. The Internet Appendix contains a description of the methodology.

Table V reports on our SVJ model calibrated to a representative Baa firm. Our approach is similar to the SV case discussed above. Our aim here is to study the joint effect of jumps and SV, together with their respective risk premia, in explaining the term structure of credit spreads. Although, as expected, including priced jumps helps fit short-term spreads and default probabilities, it does not crowd out variance risk premia, which remain important for longer maturities.

As an initial state of the world for the Baa firm, we use total asset volatility of 29%, as in Table I. We allow the jump contribution to total volatility to be

[28] Implicitly, we force all jump risk premia to be captured by the jump size premium. Like Pan (2002), we adopt this simplification out of concern for the difficulty in separately identifying both sources of risk premia. Although we do not pursue this in our estimation, the estimate for the jump size risk premium in Pan (2002) is robust to allowing for a simultaneous risk premium on the jump event timing. In a reduced-form model, where every jump leads to default, a risk premium naturally arises on the jump intensity (see Driessen (2005)). Both jump risk specifications have the advantage of generating economically significant short-term spreads.

[29] Our risk premium specification differs slightly from Pan (2002). She decomposes the total premium into premia for diffusive and jump risks, while, for expositional purposes, we prefer to have the asset drift contain a premium associated with each Brownian motion separately. Both decompositions are equivalent and imply the same shift in the volatility drift between the objective and the physical measures.

Table V

The Term Structure of Credit Spreads: Jump Risk Premia versus Variance Risk Premia

This table is based on a representative Baa firm. The exercise carried out in this table moves from a scenario with a 22% Sharpe ratio determined exclusively by variance risk premia (line 1) to scenarios with the same 22% Sharpe ratio but involving progressively greater compensation for jump risk ($\bar{\eta}_P - \bar{\eta}_Q$). The diffusive risk ($\lambda_D$) is set to zero. In each line, 1- and 10-year default probabilities are calibrated to historical levels of 0.2% and 4.9%, respectively. The jump intensity is 1%. As in the SV case, we set the correlation between the asset value and variance shocks to −0.15, the volatility of asset variance (σ) to 30%, the long-run mean variance (θ) equal to the base-case initial variance (0.29^2), and the speed of mean reversion (κ) to 4.

$\bar{\eta}_P$	$\bar{\eta}_Q$	λ_V	$QD1$	$QD10$	Spread 1Y	Spread 10Y
−0.77	−0.77	−6.81	0.0026	0.2858	25	172
−0.77	−0.84	−6.41	0.0058	0.2717	42	164
−0.77	−0.89	−6.02	0.0065	0.2611	51	157
−0.77	−0.92	−5.63	0.0071	0.2481	55	148
−0.77	−0.95	−5.24	0.0074	0.2349	56	140
−0.77	−0.98	−4.06	0.0075	0.2016	57	120

4%.[30] In all cases, we calibrate our model to a historical default probability of 0.2% for the one-year horizon and 4.9% for 10 years.[31] The mean jump size $\bar{\eta}_P$ is determined to help fit the objective default probability for a one-year horizon. The asset value drift is determined by the Sharpe ratio and total asset variance. As before, we impose an empirical Sharpe ratio of 22%. In the SVJ specification, both the jump risk premium and the volatility risk premium contribute to the level of the Sharpe ratio, which is given by

$$SR_t = \frac{\mu_X^P - r}{\sqrt{V_t}} = \frac{\left(\left(\sqrt{1-\rho^2}\lambda_D + \rho\lambda_V\right)V_t + \lambda\left(\bar{\eta}^P - \bar{\eta}^Q\right)\right)}{\sqrt{V_t + \lambda((\bar{u}^P)^2 + \gamma^2)}}. \tag{9}$$

The Sharpe ratio constraint, combined with matching historical levels for objective default probabilities, imposes a narrow range of admissible values for $\bar{\eta}_P$, $\bar{\eta}_Q$, and λ_V necessary to match the term structure of spreads. We find that the level of the jump risk premium is key to generating the one-year spread while it is mainly the variance risk premium that influences medium- to long-term maturities.

Row 1 of Table V illustrates that while a model with jumps can fit the historical default probability, it generates a relatively low one-year spread of 25 basis points. For the corporate bond data used in the calibrations, we have no observation for the one-year spread. However, for the default swap data used in our estimation, the one-year spread averages 73 basis points while the 10-year

[30] Huang and Tauchen (2005) find that about 7% of the variance in stock market index returns is due to jumps. Andersen, Bollerslev, and Diebold (2007) estimate the proportion at just less than 15%. In our case, 4% corresponds to about 13% of total asset variance.

[31] See Moody's (2010).

spread averages 130 basis points, similar to the benchmark spread of 131 basis points we use for the 10-year corporate bond. Shutting down the jump premium allows for a maximal amount of variance risk premium given the level of diffusive volatility. The implied 10-year spread, at 172 basis points, is then higher than the historical benchmark. As we progressively allow for a greater jump risk premium, the complementarity of the jump and volatility risk premia in matching different parts of the term structure becomes apparent. Keeping $\bar{\eta}_P$ constant to match the objective default probabilities, the level of $\bar{\eta}_Q$ controls the relative importance of the jump risk premium. Given our Sharpe ratio constraint, the increase in jump risk premium comes at the expense of the variance risk premium.

As we increase the jump risk component, one-year spreads increase rapidly at first and then taper off. This is due to the fact that there is some level of $\bar{\eta}_Q$ that is sufficiently large to trigger default. Beyond this level, increasing $\bar{\eta}_Q$ has no significant effect.[32] However, as jump risk premia and short-term spreads increase, λ_V is reduced and this induces lower 10-year spreads. Between the last two rows, there is a combination of jump and variance risk premia that produces the benchmark 10-year historical spread of 131 basis points with one-year spreads more than twice as high as in the no jump risk premium case.

Note that, although the jump risk premium to some degree "crowds out" the variance risk premium, the effect is not strong enough to preclude the premium for variance risk from playing an important role. The level of λ_V that together with jump risk matches 10-year spreads is slightly higher than what we found in Table I. Note, however, that this higher level of λ_V is applied to a lower level of diffusive volatility (25% versus 29% in Table I).

In summary, given a reasonable Sharpe ratio, our model without jumps cannot generate sufficient spread levels or short-term default probabilities. Jumps help fit the short-term default probability, while the risk premium helps with the short-term spread. The presence of jumps somewhat reduces the amount of permissible variance risk, but not enough to preclude it from helping explain medium- and long-term spreads as before.

V. Estimating the SVJ Model

A. *Estimation Methodology*

For a given firm, we have seven observables: the CDS spreads at 1, 2, 3, 5, 7, and 10 years plus the equity volatility. Since we have only two states, namely, asset value and variance, we follow the usual practice in the literature (e.g., Duffee (2002)) and assume that two variables, jointly denoted by $Y_t^{(a)}$, are accurately observed, while the other five variables, jointly denoted by $Y_t^{(m)}$, are observed with error. Specifically, we choose $Y_t^{(a)} = \{cds_t(5), \sigma_{E,t}\}$ and $Y_t^{(m)} = \{cds_t(1), cds_t(2), cds_t(3), cds_t(7), cds_t(10)\}$. Our choice reflects the stance

[32] This limitation is an artifact of an exogenous financial distress cost proportional to the default boundary. In a more elaborate model that may consider distress costs as a function of the jump severity, an increase in the size of the jump may continue to contribute to larger spreads.

that (i) five-year contracts are the most liquid and (ii) $\sigma_{E,t}$ can be accurately measured using option-implied volatilities. For notational convenience, we denote by $Y_t = \{Y_t^{(a)}, Y_t^{(m)}\}$.

The parameters η, α, and L, which denote the effective corporate tax rate, the firm-level distress cost, and the bond-specific default loss, are set to 15%, 15%, and 49%, respectively, consistent with prior empirical work. For a given firm, we obtain the time series of its payout ratios and debt service rates, and use their averages as calibrated values for q and c. Denote by $\Theta = \{\lambda_D, \lambda_V, \kappa, \theta, \sigma, \rho, X_d, \bar{\eta}^Q, \lambda\}$ the remaining parameters to be estimated, where λ_D determines the diffusive risk premium; parameters κ, θ, and σ characterize the V_t dynamics; λ_V determines (κ^*, θ^*) (the risk-adjusted counterparts of (κ, θ)); ρ denotes the correlation between firm asset value and variance shocks; X_d is the (constant) default boundary; $\bar{\eta}^Q$ is the average jump size under Q; and λ is the jump intensity under Q.

To identify the parameter vector Θ, we use ML estimation with respect to the P-distribution of the observables Y_t. Note that Θ contains some parameters relating to the Q dynamics of the model. ML estimation requires the computation of CDS spreads, which helps identify the Q-measure parameters. In addition, the derivation of the P-distribution of Y_t depends on the P-distribution of the two states (z_t, V_t), which helps us identify its parameters.[33]

Define $S_t \equiv \{z_t, V_t\}$, the two estimated states in our model, where $z_t \equiv ln(X_t/X_D)$. Given the parameter vector Θ, the accurately observed data $Y_t^{(a)}$ can be inverted to form an implied state vector $\hat{S}_t$. To facilitate the subsequent exposition, we write

$$\hat{S}_t = F\left(Y_t^{(a)}\right), \tag{10}$$

where $F(.)$ denotes the functional form that maps $Y_t^{(a)}$, the two variables assumed to be accurately observed, into the two states. To see the identification of S_t from $Y_t^{(a)}$, first note that $cds_t(5) = Y_t^{(a)}(1)$ is monotone in $z_t = S_t(1)$—intuitively, a higher $z_t(\equiv ln(X_t/X_D))$ implies longer distance to default and lower credit spreads. Second, $\sigma_{E,t} = Y_t^{(a)}(2)$ is monotone in $V_t = S_t(2)$, since a higher asset value variance is associated with a higher equity variance. Numerically, we back out $\hat{S}_t$ from $Y_t^{(a)}$ using a two-stage grid search procedure that works uniformly well for all firms.[34]

Given $\hat{S}_t$, the model-implied spreads for the other five CDS contracts, which we denote by $\hat{Y}_t^{(m)}$, can be calculated. Let $\epsilon_t^R \equiv \frac{\hat{Y}_t^{(m)} - Y_t^{(m)}}{Y_t^{(m)}} \in \mathbb{R}^5$ be the relative

[33] In an earlier version of this paper, we use GMM estimation instead. An important caveat of GMM is that the model-implied pricing moments are determined only by Q-measure parameters, which leaves P-measure parameters unidentifiable.

[34] More specifically, we choose a wide range of z and V in the first stage that cover all firms. At the end of the first stage, we identify, for each firm, a subrange in which the implied states are most likely to reside. We then conduct the second-stage finer grid search that precisely defines the implied $\hat{S}_t$. In both stages, we use function approximation by means of Chebychev polynomials to represent the dependence of $Y_t^{(a)}$ on the states. This greatly improves the computation speed while maintaining accuracy.

observation error. We assume that the variance-covariance matrix of ϵ_t^R is time invariant, which implies the Cholesky decomposition $E(\epsilon_t^R(\epsilon_t^R)') = CC'$. For simplicity, we choose $C = \sigma_e I_{5\times 5}$, where σ_e denotes the standard deviation of the measurement error and $I_{5\times 5}$ denotes the 5-by-5 identity matrix, and conduct estimations for all firms in our sample under this specification.[35]

What remains is to derive p_Y, the P-distribution of $Y_t^{(a)}$. In our model, the P-measure state transition density of S_T conditional on S_t is not available in closed form. We therefore use the simulated maximum likelihood estimation (SMLE) method proposed by Brandt and Santa-Clara (2002) and resort to simulation to approximate the state transition density $p(S_{t_{n+1}}, t_{n+1}|S_{t_n}, t_n)$ for two adjacent discrete observation times t_n and t_{n+1}. We then apply the mapping from our states to the implied $Y_t^{(a)}$, as summarized by (10), to obtain p_Y from p.

Normalizing the length of $[t_n, t_{n+1}]$ to one, we choose the popular Euler discretization scheme and divide this interval into M subintervals of length $h = 1/M$. The Euler discretization of the two state vectors $S_t = (z_t, V_t)$ in $[t_n, t_{n+1}]$, denoted by $(\hat{z}_{t_n+mh}, \hat{V}_{t_n+mh})$ for $m = 0, 1, \ldots, M-1$, follows

$$\begin{bmatrix} \hat{z}_{t_n+(m+1)h} \\ \hat{V}_{t_n+(m+1)h} \end{bmatrix} = \begin{bmatrix} \hat{z}_{t_n+mh} \\ \hat{V}_{t_n+mh} \end{bmatrix} + \nu(.)h + \Sigma(.)\sqrt{h}\begin{bmatrix} \varepsilon^z \\ \varepsilon^V \end{bmatrix} + \begin{bmatrix} U_z \\ 0 \end{bmatrix} N(h), \tag{11}$$

where

$$\nu(.) = \begin{bmatrix} r - q + \left(\sqrt{1-\rho^2}\lambda_D + \rho\lambda_V - \frac{1}{2}\right)\hat{V}_{t_n+mh} - \lambda^Q\bar{\eta}^Q \\ \kappa(\theta - \hat{V}_{t_n+mh}) \end{bmatrix}, \tag{12}$$

$$\Sigma(.) = \begin{bmatrix} \sqrt{1-\rho^2}\sqrt{\hat{V}_{t_n+mh}} & \rho\sqrt{\hat{V}_{t_n+mh}} \\ 0 & \sigma\sqrt{\hat{V}_{t_n+mh}} \end{bmatrix}, \tag{13}$$

$[\varepsilon^z, \varepsilon^V]'$ is the two-dimensional standard normal random variates, $U_z \sim^P N(\bar{u}^P, \gamma^2)$, and $N(h)$ denotes the Poisson process over a small time interval h with intensity λh.

Denote by $\psi(S_{t_{n+1}}|S_{t_n+(M-1)h})$ the one-step-ahead state transition density of the Euler discretization from $t_n + (M-1)h$ to $t_{n+1}(= t_n + Mh)$. The original SMLE as in Brandt and Santa-Clara (2002) relies on pure diffusions, under which ψ can simply be approximated by (joint) normal distributions. As this is not the case with jumps, we simplify our analysis by assuming that there is at most one jump over a small time interval h. More specifically, if there is no jump between $t_n + (M-1)h$ and t_{n+1}, ψ takes the approximation form

$$\phi^{(2)}\left(\begin{bmatrix} \hat{z}_{t_{n+1}} \\ \hat{V}_{t_{n+1}} \end{bmatrix}; \begin{bmatrix} \hat{z}_{t_n+(M-1)h} \\ \hat{V}_{t_n+(M-1)h} \end{bmatrix} + \nu(.)h,\ \Sigma(.)\Sigma(.)'h\right), \tag{14}$$

[35] We have considered alternative specifications for C, and the results continue to hold.

where $\phi^{(2)}$ denotes a bivariate normal density. If a jump does occur, its impact on the z-innovation between $t_n + (M-1)h$ and t_{n+1} dominates that from the diffusions over a small interval h. We can thus attribute all of $\hat{z}_{t_n+(M-1)h} - \hat{z}_{t_{n+1}}$ to the realized jump, and in this case ψ can be approximated by

$$\phi^{(1)}\Big(\hat{V}_{t_{n+1}}; \hat{V}_{t_n+(M-1)h} + \kappa(\theta - \hat{V}_{t_n+(M-1)h})h, \sigma^2 \hat{V}_{t_n+(M-1)h} h\Big) \times \phi^{(1)}\Big(\hat{z}_{t_{n+1}}; \bar{u}^P, \gamma^2\Big), \quad (15)$$

where $\phi^{(1)}$ denotes a univariate normal density. Specifically, (15) follows from (i) the fact that the V-diffusions are independent of z-jumps and (ii) the jump size of z, U_z, is normally distributed with mean $\bar{u}^P$ and variance γ^2. The approximations of both (14) and (15) are exact as $h \to 0$.

In view of the Euler discretization and the definition of ψ, the state transition density $p(S_{t_{n+1}}, t_{n+1}|S_{t_n}, t_n)$ can be approximated through recursive integration as follows:

$$p^M\big(S_{t_{n+1}}, t_{n+1}|S_{t_n}, t_n\big) = \int\int \psi\big(S_{t_{n+1}}|x\big) \times p^M\big(x, t_n + (M-1)h|S_{t_n}, t_n\big)dx, \quad (16)$$

where x denotes a particular realization of $S_{t_n+(M-1)h} \in R \times R_+$. We can interpret the integral in (16) as the expectation of the function ψ of the random variable x, where the distribution of x is $f(x) = p^M(x, t_n + (M-1)h|S_{t_n}, t_n)$. Again following Brandt and Santa-Clara (2002), we use Euler discretization as in (11) to simulate a large number of independent random draws x_j for $j = 1, 2, \ldots, N$ from the distribution $f(x)$.[36] We then approximate the expectation, and ultimately the corresponding transition density p, with a sample average of the function ψ evaluated at these random draws of x, that is,

$$\hat{p}^{M,N}\big(S_{t_{n+1}}, t_{n+1}|S_{t_n}, t_n\big) = \frac{1}{N}\sum_{j=1}^{N} \psi\big(S_{t_{n+1}}|x_j\big)dx, \quad (17)$$

where x_j for $j = 1, 2, \ldots, N$ represents independent realizations of $S_{t_n+(M-1)h}$ after $M-1$ Euler recursions according to (11) starting from S_{t_n}. The Strong Law of Large Numbers guarantees that $\hat{p}^{M,N}$ converges to the transition density p as $N, M \to \infty$.[37]

To implement (17), we need to decide, for a given path of the simulated $x = (\hat{z}_{t_n+(M-1)h}, \hat{V}_{t_n+(M-1)h})$, on the threshold under which the z-innovation between $t_n + (M-1)h$ and t_{n+1} is attributed to a jump. We pick this threshold to be $\frac{1}{4}\bar{u}^P (< 0)$, where $\bar{u}^P$ is the average jump size of z_t under P. Our choice accounts for the stochastic nature of the jump magnitude, and it roughly matches the jump intensity of λh between $t_n + (M-1)h$ and t_{n+1}. Thus, ψ takes (14) and (15) for $\hat{z}_{t_{n+1}} - \hat{z}_{t_n+(M-1)h} >$ and $< \frac{1}{4}\bar{u}^P$, respectively. Under this scheme and at

[36] More precisely, for each x_s, we start at time t_n and iterate through the Euler recursion (11) exactly $M-1$ times to get one x_s. We repeat this procedure S times, which yields the random sample $\{x_1, x_2, \ldots, x_S\}$.

[37] In more detail, $\hat{p}^{M,N}$ converges to p^M as $N \to \infty$, and p^M converges to p as $M \to \infty$ (or equivalently $h \to 0$). See Brandt and Santa-Clara (2002) for the detailed proofs.

the monthly frequency of the observed data, we find that an M of 30 to 40 and an N of 50,000 to 100,000 are sufficient to achieve convergence.

Given the simulated $p\left(S_{t_{n+1}}, t_{n+1}|S_{t_n}, t_n\right)$, the P-distribution of $Y^{(a)}_{t_{n+1}}$ conditional on $Y^{(a)}_{t_n}$ is given by

$$p_Y\left(Y^{(a)}_{t_{n+1}}, t_{n+1}|Y^{(a)}_{t_n}, t_n\right) = \frac{1}{\left|\det\left(J_{t_{n+1}}\right)\right|} p\left(S_{t_{n+1}}, t_{n+1}|S_{t_n}, t_n\right), \tag{18}$$

where

$$J_t \equiv \begin{bmatrix} \frac{\partial Y^{(a)}_t(1)}{\partial S_t(1)} & \frac{\partial Y^{(a)}_t(1)}{\partial S_t(2)} \\ \frac{\partial Y^{(a)}_t(2)}{\partial S_t(1)} & \frac{\partial Y^{(a)}_t(2)}{\partial S_t(2)} \end{bmatrix} = \begin{bmatrix} \frac{\partial cds_t(5)}{\partial z_t} & \frac{\partial cds_t(5)}{\partial V_t} \\ \frac{\partial \sigma_{E,t}}{\partial z_t} & \frac{\partial \sigma_{E,t}}{\partial V_t} \end{bmatrix}. \tag{19}$$

In (19), the relation between $Y^{(a)}_t$ and S_t is due to the implied mapping in (10). Specifically, we use Chebychev interpolation (see, e.g., Miranda and Fackler (2002)) to fit model spreads as functions of the states, which also significantly facilitates the computation of J_t.

Returning to $Y^{(m)}_t$, we denote their joint distribution by $f_\epsilon(\epsilon_t; Y^{(m)}_t)$, where $\epsilon_t \in \mathbb{R}^5$ is the relative observation error, which is assumed to be normally distributed with variance-covariance matrix $\sigma_e^2 I_{5\times 5}$. The log likelihood function for observations at t_n is then

$$\begin{aligned} l_{t_n}(\Theta) &= \log p_Y\left(Y^{(a)}_{t_n}, t_n|Y^{(a)}_{t_{n-1}}, t_{n-1}\right) + \log f_\epsilon\left(\epsilon_{t_n}; Y^{(m)}_{t_n}\right) \\ &= \log p\left(S_{t_n}, t_n|S_{t_{n-1}}, t_{n-1}\right) - \log\left|\det\left(J_{t_n}\right)\right| + \log f_\epsilon\left(\epsilon_{t_n}; Y^{(m)}_{t_n}\right), \end{aligned} \tag{20}$$

where $p(S_{t_1}, t_1|S_{t_0}, t_0)$ denotes the simulated unconditional state distribution. Due to a relatively long time series of our sample, we find that assuming a degenerate $p(S_{t_1}, t_1|S_{t_0}, t_0)$ incurs very limited information loss compared to that using the unconditional distribution of S_t. Our estimation is thus based on

$$\max_{\Theta} L(\Theta) = \sum_{n=2}^{T} l_{t_n}(\Theta), \tag{21}$$

where $\Theta = \{\lambda_D, \lambda_V, \kappa, \theta, \sigma, \rho, X_d, \bar{\eta}^P, \bar{\eta}^Q, \sigma, \lambda\}$, and T denotes the number of observations. Specifically, we use monthly data, with T equal to 84 for most firms.

We consider various ways to choose reasonable initial values for Θ. For the two risk premium parameters, we perform a calibration similar to Table I for each rating category and use the implied λ_D and λ_V as their initial values. For the rest of the nonjump parameters, $\{\kappa, \theta, \sigma, \rho, X_d\}$, we set their (firm-by-firm) initial values to their GMM estimates without jumps from a previous version of our paper. For the jump-related parameters, we set the initial intensity value based on short-term CDS spreads normalized by the loss rate. The jump sizes are chosen based on short-term historical default rates by credit rating.

B. Data

We collect single-name CDS spreads from Markit. Daily CDS spreads reflect the average quotes contributed by major market participants. This database has already been cleaned to remove outliers and stale quotes. We require that two or more banks contribute spread quotes to include an observation (Cao, Yu, and Zhong (2010)). The data sample available to us includes only those firms that constituted the CDX index from January 2001 to December 2013. Our sample includes U.S. dollar-denominated CDS contracts written on senior unsecured debt of U.S. firms. We use maturities of 1, 2, 3, 5, 7, and 10 years. Table VI summarizes our CDS data and key firm-level variables.

The range of restructurings that qualify as credit events varies across CDS contracts from no restructuring (XR) to unrestricted restructuring (CR). Modified restructuring (MR) contracts that limit the range of maturities of deliverable instruments in the case of a credit event are the most popular contracts in the United States. We therefore include only U.S. dollar–denominated contracts on senior unsecured obligations with modified restructuring (MR).

Following Huang and Zhou (2008), we perform our test on monthly data. Their sample is restricted to 36 monthly intervals as their sample ends in 2004. Instead, we require that the CDS time series have at least five years of consecutive monthly observations to be included in the final sample. We also require that CDS data have matching equity prices (CRSP), equity volatility (TAQ), and accounting variables (Compustat). We further exclude the financial and utility sectors, following previous empirical studies on structural models. After applying these filters, we are left with 49 entities in our study.

In testing structural models, asset return volatility is unobserved and is often backed out from the observed equity return volatility. We collect implied volatility data from the OptionMetrics database on the last trading day of each month. Commonly used option data filters including those proposed by Bakshi, Cao, and Chen (1997) are applied. Specifically, we discard all options with zero volume, zero open interest, and best bid price less than \$0.1, as well as those that violate no-arbitrage bounds. We then pick the at-the-money (ATM) options with the shortest time to maturity. We define ATM by its moneyness using the Black-Scholes delta closest to 0.5 for calls and -0.5 for puts.

C. Estimation Results

In this section, we discuss our findings from a firm-by-firm estimation of the SVJ model. Although we proceed directly to the SVJ case in our empirical exercise, we also carried out an SV-only estimation, the results of which are reported in the Internet Appendix. Table VII provides a summary of the parameter estimates.[38]

Perhaps the most important parameter for our purposes is λ_V. We obtain a broad range of estimates with 95% in the $[-2.80, -0.01]$ range. The mean

[38] Firm-by-firm parameter estimates and standard errors for the SVJ estimation are provided in the Internet Appendix.

Table VI

CDS and Firm Data

This table provides details on the term structure of CDS spreads across maturities of 1, 2, 3, 5, 7, and 10 years. For each firm and maturity, we report mean spreads over the time series for each firm, as well as the standard deviation along the same dimension. We also report mean leverage ratios, computed as book debt over book debt plus market capitalization; equity volatility, computed using intraday data; and the asset payout rate, computed as the weighted average of dividend payout and interest expenses relative to total assets.

		One-Year		Two-Year		Three-Year		Five-Year		Seven-Year		10-Year		Leverage	Equity	Payout	Start	End
Firm	Credit rating	Mean	Std. Dev	Mean	Std. Dev	Mean	Std. Dev	Mean	Std. Dev	Mean	Std. Dev	Mean	Std. Dev	Ratio	Volatility	Rate	Date	Date
Amgen	A	15	7	24	8	33	9	54	13	68	13	82	15	0.31	0.61	0.018	Apr-10	Dec-13
Arrow Electronics	BBB	90	113	97	103	116	102	147	93	161	89	172	83	0.58	0.37	0.016	Mar-01	Dec-13
AT &T	A	64	124	73	121	80	117	93	111	100	106	109	102	0.54	0.24	0.035	Jan-01	Dec-13
Baxter International	A	18	18	22	17	26	17	35	15	41	14	48	14	0.26	0.25	0.017	Apr-01	Dec-13
Boeing	A	30	36	37	37	43	38	57	42	64	41	73	41	0.53	0.29	0.017	Mar-02	Dec-13
Boston Scientific	BB	49	54	62	58	77	62	100	66	114	65	120	62	0.34	0.4	0.01	Apr-01	Dec-13
Bristol Myers Squibb	A	14	12	18	12	23	12	32	13	39	13	46	14	0.25	0.26	0.037	Jan-01	May-13
Campbell Soup	A	15	9	20	9	26	12	38	17	46	22	54	26	0.33	0.21	0.028	Jan-01	Dec-13
Caterpillar	A	32	46	39	46	46	47	61	51	69	51	77	51	0.55	0.31	0.024	Jul-02	Dec-13
Cigna	BBB	44	54	53	55	63	56	81	57	91	55	100	53	0.8	0.35	0.003	Nov-02	Dec-13
Comcast	BBB	43	56	55	60	66	59	88	57	100	53	112	48	0.62	0.3	0.022	Jan-01	Jul-13

(*Continued*)

Table VI—*Continued*

Firm	Credit rating	One-Year Mean	One-Year Std. Dev	Two-Year Mean	Two-Year Std. Dev	Three-Year Mean	Three-Year Std. Dev	Five-Year Mean	Five-Year Std. Dev	Seven-Year Mean	Seven-Year Std. Dev	10-Year Mean	10-Year Std. Dev	Leverage Ratio	Equity Volatility	Payout Rate	Start Date	End Date
Computer Sciences	BBB	44	58	57	65	71	73	103	92	118	99	132	102	0.51	0.35	0.013	Sep-01	Mar-13
ConAgra Foods	BBB	23	15	31	15	39	18	56	25	66	31	77	34	0.42	0.21	0.037	Feb-01	Dec-13
Darden Restaurants	BBB	46	54	57	54	70	57	96	65	110	68	122	69	0.34	0.33	0.019	Jan-01	Oct-13
Devon Energy	BBB	36	40	43	38	52	39	67	39	76	40	86	41	0.43	0.34	0.018	Jan-01	Sep-13
Disney	A	24	28	30	30	35	28	46	27	54	24	65	23	0.33	0.29	0.015	Jan-01	Dec-13
Dow Chemical	A	59	94	69	94	79	93	100	90	110	86	120	84	0.5	0.33	0.034	Dec-01	May-13
Dupont	A	18	19	23	22	29	23	41	27	49	28	57	29	0.4	0.26	0.028	Feb-01	Dec-13
Eastman Chemical	BBB	36	29	44	28	55	28	78	30	91	32	104	33	0.5	0.3	0.028	Jan-01	Dec-11
Eastman Kodak	B	590	2,013	592	1,584	607	1,402	634	1,227	617	1,114	601	1,006	0.67	0.49	0.022	Jul-01	Dec-13
Ford Motor	B	654	1,735	699	1,580	722	1,447	738	1,292	721	1,194	699	1,092	0.88	0.44	0.03	Feb-01	Dec-13
Fortune Brands	BBB	42	48	52	50	61	53	79	59	89	56	99	56	0.4	0.26	0.024	Jan-01	Sep-11
General Electric	AAA	73	144	82	138	88	133	100	129	103	118	107	113	0.67	0.28	0.028	Jan-01	Dec-13
General Mills	BBB	24	21	29	20	35	20	46	18	54	18	63	18	0.39	0.19	0.03	Apr-01	May-13
Goodrich	BBB	49	60	53	57	59	57	73	56	81	54	89	54	0.49	0.33	0.023	Mar-01	Mar-12

(*Continued*)

Table VI—*Continued*

Firm	Credit rating	One-Year Mean	One-Year Std. Dev	Two-Year Mean	Two-Year Std. Dev	Three-Year Mean	Three-Year Std. Dev	Five-Year Mean	Five-Year Std. Dev	Seven-Year Mean	Seven-Year Std. Dev	10-Year Mean	10-Year Std. Dev	Leverage Ratio	Equity Volatility	Payout Rate	Start Date	End Date
Honeywell	A	20	21	25	22	29	21	38	22	45	21	53	20	0.4	0.28	0.021	Jul-01	Dec-13
Ingersoll Rand	BBB	47	71	63	69	80	86	106	98	112	87	121	87	0.47	0.34	0.018	Jan-01	Dec-13
International Paper	BBB	70	110	84	113	96	110	121	105	134	98	147	92	0.59	0.33	0.031	Feb-06	Dec-13
JC Penney	A	134	101	148	92	171	90	197	80	208	77	210	71	0.5	0.25	0.012	Jan-01	Jan-13
John Deere	A	25	27	31	29	37	30	49	31	56	29	64	29	0.6	0.32	0.025	Oct-01	Sep-12
Kraft Foods	BBB	26	23	32	23	39	24	53	25	61	25	71	26	0.56	0.21	0.038	Jan-01	Jan-13
Kroger	BBB	37	27	46	25	56	26	77	27	88	29	99	31	0.53	0.28	0.02	Apr-01	Apr-06
Marriott International	BBB	63	91	74	94	84	95	104	94	114	86	125	82	0.35	0.33	0.015	Jan-01	Dec-13
McDonald's	A	13	11	17	11	21	11	30	12	36	11	43	11	0.23	0.23	0.025	Jan-01	Dec-13
McKesson	BBB	32	32	37	32	45	33	59	33	68	34	79	38	0.54	0.28	0.009	Jan-01	Mar-13
Newell Rubbermaid	BBB	41	40	50	41	61	43	84	47	96	48	108	48	0.44	0.3	0.028	Feb-01	Mar-11
Northrop Grumman	BBB	30	28	34	27	42	31	53	32	59	31	65	31	0.48	0.24	0.021	Jan-01	Dec-13
Olin	BB	69	66	78	56	100	57	136	54	149	52	162	53	0.48	0.34	0.03	Jan-02	Dec-13
Radioshack	BB	241	385	299	465	335	502	369	526	376	518	389	502	0.4	0.48	0.022	Jun-01	Dec-13

(*Continued*)

Table VI—*Continued*

Firm	Credit rating	One-Year Mean	One-Year Std. Dev	Two-Year Mean	Two-Year Std. Dev	Three-Year Mean	Three-Year Std. Dev	Five-Year Mean	Five-Year Std. Dev	Seven-Year Mean	Seven-Year Std. Dev	10-Year Mean	10-Year Std. Dev	Leverage Ratio	Equity Volatility	Payout Rate	Start Date	End Date
Raytheon	BBB	37	46	41	42	49	45	61	42	69	41	77	40	0.45	0.24	0.023	May-01	Dec-13
Safeway	BBB	39	25	53	36	69	52	104	84	121	99	135	105	0.52	0.31	0.023	Apr-01	Dec-13
Sherwin-Williams	A	26	25	32	24	40	25	55	27	64	27	73	30	0.3	0.29	0.017	Feb-01	Feb-13
SuperValu	BB	213	204	258	243	290	266	339	298	353	295	364	282	0.7	0.39	0.038	Feb-01	Dec-13
Target	A	26	32	30	34	36	36	46	37	52	35	61	34	0.38	0.31	0.017	Jan-01	Dec-13
Temple Inland	BB	110	155	123	151	136	148	161	139	171	128	183	116	0.79	0.4	0.037	Jan-01	Jan-13
Tyson Foods	BB	81	85	97	85	116	85	149	81	164	76	177	72	0.59	0.3	0.03	Apr-01	Sep-11
Verizon Communications	A	71	104	74	95	77	88	82	71	87	68	97	61	0.5	0.27	0.034	Aug-01	Sep-13
Walmart	AA	16	16	20	18	24	19	31	21	36	21	42	22	0.29	0.23	0.016	Jan-01	Jan-06
Whirlpool	BBB	58	83	71	84	85	85	114	90	127	88	139	86	0.61	0.37	0.021	Jan-01	Dec-13
Mean		75		85		96		115		124		158		0.48	0.262	0.0234		
Median		41		52		61		79		89		81		0.49	0.261	0.0227		

Table VII
Summary Statistics of SVJ Parameter Estimates

This table summarizes the parameter estimates obtained for our estimation of the SVJ model. We denote the variance risk premium parameter by λ_V, the speed of mean reversion by κ, the long-run mean variance by θ, the volatility of variance by σ, the correlation between the Brownian motions defining asset value and variance shocks by ρ, and the spread observation error by σ_e. The jump parameters are the risk-adjusted mean jump size, $\bar{\eta}_Q$, and the jump intensity, λ. The last two columns provide the risk-adjusted version of the speed and level of mean reversion, κ^* and θ^*. Panel B reports the corresponding standard errors.

Panel A: Parameter estimate quantiles

Quantile	λ_V	κ	θ	σ	ρ	$\bar{\eta}_Q$	λ	σ_e	κ^*	θ^*
5	−2.796	0.463	0.009	0.110	−0.296	−0.938	0.001	0.197	0.155	0.027
25	−2.281	1.926	0.017	0.305	−0.215	−0.901	0.002	0.288	1.229	0.027
50	−1.756	3.231	0.030	0.398	−0.157	−0.803	0.004	0.341	2.533	0.038
75	−1.134	4.028	0.042	0.610	−0.116	−0.650	0.009	0.416	3.336	0.050
95	−0.098	5.965	0.070	0.835	−0.006	−0.295	0.046	0.782	5.883	0.071
Mean	−1.675	3.146	0.032	0.447	−0.161	−0.810	0.010	0.392	2.397	0.042

Panel B: Standard error quantiles

Quantile	λ_V	κ	θ	σ	ρ	$\bar{\eta}_Q$	λ	σ_e
5	0.0181	0.0347	0.0005	0.0070	0.0017	0.0159	0.0000	0.0211
25	0.0520	0.1713	0.0019	0.0159	0.0058	0.0692	0.0002	0.0326
50	0.1130	0.3570	0.0036	0.0372	0.0148	0.1223	0.0003	0.0434
75	0.2722	0.7760	0.0065	0.0821	0.0310	0.3706	0.0008	0.0764
95	1.2714	3.6925	0.0351	0.5721	0.1542	2.0774	0.0180	0.2862
Mean	0.3601	0.9577	0.0078	0.1374	0.0400	0.5153	0.0048	0.1095

is −1.675, which is smaller in absolute terms compared to what we found in our calibrations. The mean reversion speed parameter κ averages 3.15, which is in line with estimates from S&P 500 returns in, for example, Bates (2006). The average of the long-run volatility (θ) is 0.032. The estimation error averages 0.392. We also report risk-adjusted versions of the speed and level of mean reversion.

The results in Table VII confirm that it is large infrequent jumps that the data requires. The average intensity across firms is around 0.5%, with cross-sectional variation from 0.1% to 4.6%. The firm-by-firm average one-year spread in our data varies from 13 basis points to 654 basis points.[39] The mean Q jump size $\bar{\eta}^Q$ is significant across all firms with an average of −81%. It is interesting to note that we do have firms for which the jump intensity and size are zero. These are cases in which we are unable to identify jump risk parameters that improve the likelihood over the best SV estimates. Most of these firms also have well below average one-year spreads. For those firms, it appears that variance risk is enough to explain the term structure of spreads.

[39] The one-year spread for Ford Motor reached 1,222 basis points in February 2009.

Table VIII
Pricing Performance of the SVJ Model

This table reports summary statistics for root mean squared pricing errors (RMSE) in basis points, as well as relative RMSE (RRMSE), resulting from our estimation of the SVJ model on 49 firms between 2001 and 2013.

Quantile	RMSE 1	RMSE 2	RMSE 3	RMSE 7	RMSE 10
5	9	6	4	5	8
25	20	11	8	7	12
50	27	20	13	10	18
75	56	36	23	15	25
95	438	227	154	65	114
Mean	75	44	32	25	36
Quantile	RRMSE 1	RRMSE 2	RRMSE 3	RRMSE 7	RRMSE 10
5	0.40	0.24	0.16	0.08	0.13
25	0.46	0.29	0.19	0.10	0.15
50	0.52	0.32	0.22	0.12	0.18
75	0.69	0.42	0.26	0.14	0.22
95	1.00	0.59	0.37	0.21	0.30
Mean	0.61	0.37	0.24	0.13	0.19

The variation in $\bar{\eta}^Q$ demonstrates that the jump risk premium varies across firms.[40] Note that the jump risk premium, defined as the ratio between risk-adjusted to physical jump sizes, is smaller in our analysis compared to Pan (2002).[41] Pan (2002) examines the jump premium embedded in an equity index, while we focus on firm-by-firm estimation of asset value jumps. First, jumps in equity are amplified by leverage. Second, diversification in the index would lead to the remaining jumps being systematic and priced. At the firm level, on the other hand, some of the jump risk is idiosyncratic and does not carry a premium.

Turning our attention to the in-sample fit, we present the RMSE for maturities of 1, 2, 3, 7, and 10 years. We omit the five-year maturity and equity volatility because they are used to filter the states and hence exhibit no pricing errors. Summaries of the errors are reported in Table VIII while firm-by-firm results are reported in the Internet Appendix. The largest pricing errors are to be found at the shorter end of the maturity spectrum. There is no particular pattern across absolute and relative errors in the cross-section. In the time-series

[40] Elkamhi and Ornthanalai (2010) document the presence and cross-sectional variation of equity jump risk and jump risk premia at the firm level.

[41] Pan (2002) estimates a risk-adjusted jump size of about −19% with a physical counterpart of about −1%. While our mean physical and risk-adjusted jump size are much larger at −70% and −81%, respectively, the differences and ratios of the two kinds of jump sizes are larger in Pan (2002).

dimension, however, basis point (and to a lesser extent proportional errors) increase significantly during the recent financial crisis.[42]

Overall, we find that in an SVJ model, the variance risk and jump risk premia are both important. Figure 3 compares the pricing performance of our SVJ model with two nested models: a constant volatility (CV) model comparable to Black and Cox (1976) and Leland (1994), and our model with jumps removed (SV). In terms of predicting the level of spreads, we see that the gains from the SV specification over the CV model are particularly significant in the 1 to 3 years to maturity bracket. The bias at one year is halved and at two years almost entirely gone. The additional gains from the inclusion of jumps in this aspect is an additional significant reduction in bias for the shortest, one-year maturity. In terms of root mean squared pricing errors, a similar picture emerges. The biggest improvement takes place when moving from the CV to SV model, particularly at the shorter maturities. The most significant gains from the inclusion of jumps are also made at the short end.

While we show in Section IV that our simple jump structure helps generate the levels of short-term spreads, the story is more nuanced when the model is estimated at the firm-by-firm level. In the estimation, the choice of securities to fit, as well as the loss function, may implicitly favor diffusive risk over jumps. As we have seen in the calibration section, medium and long maturities are more sensitive to volatility risk. There is only one short maturity (one year) in our data.[43] The other CDS maturities used in our estimation are 2, 3, 5, 7, and 10. Thus, by construction, our likelihood places more weight on medium- to long-term maturities, which are also more sensitive to variance risk than jump risk. Matching shorter maturities requires jumps. But as we have seen, increasing short-term spreads by means of more jump risk comes at the expense of the ability to generate long-term spreads (since we match total equity volatility). Hence, our SMLE procedure may sacrifice the ability to fit short-term spreads in order to fit the medium to long term. In addition, the tension between short- and long-term spreads is exacerbated in periods of upward-sloping spread term structures, which are common to most of our data save during the financial crisis.

Another factor in the underestimation of short-term spreads is that the one-year spread may be more influenced by illiquidity risk than other maturities. Since our model abstracts from illiquidity risk, a downward bias in short-term spreads may in fact be reasonable.[44]

An advantage of our structural credit risk model is that, as a by-product of our estimation, we can compute other firm-level quantities. For example, since we identify both physical and risk-adjusted asset value and volatility dynamics,

[42] There is evidence that fixed income markets experienced significant disruptions during this period: Bid-ask spreads were high (Marra (2017)), and the CDS-bond basis diverged and became negative (Bai and Collin-Dufresne (2018)), suggesting that arbitrage activity between markets became more difficult. This may explain the larger pricing errors for our model during this period.

[43] Note that the maturities considered in the large literature on SV in equity options markets are much shorter.

[44] We consider an extension to address illiquidity risk in the Internet Appendix.

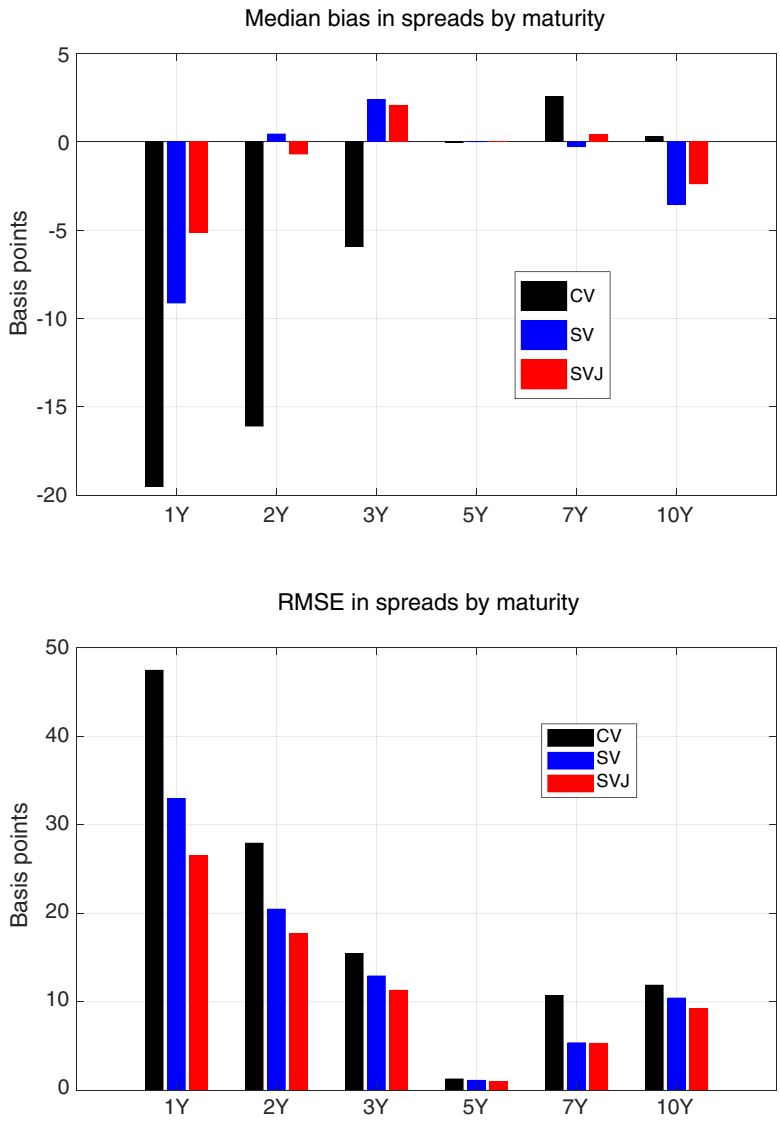

Figure 3. Relative performance of the SVJ and the nested SV and CV models. This figure summarizes the pricing performance of the SVJ model specification and compares it to the nested stochastic volatility (SV) and constant volatility (CV) models. The top panel in this figure plots the median bias, that is, the model minus market spreads in basis points, per maturity segment. The median bias is calculated across the predicted spreads for the 49 firms in our sample. The second panel plots the basis point root mean squared error for the two model specifications across maturity segments for the same sample of firms. Recall that, in our estimation, the five-year CDS spread is assumed to be perfectly observed, so no pricing error results. (Color figure can be viewed at wileyonlinelibrary.com)

Table IX
Predicted Economic Quantities

This table reports summary statistics for various model-predicted economic quantities as well as pricing errors resulting from our estimation of the SVJ model for 49 firms between 2001 and 2013. The equity variance risk premium (EVRP) is computed using equation (5). The Sharpe ratio is computed using equation (1).

		Physical Default Prob.			Sharpe
Quantile	EVRP	One-Year	Four-Year	10-Year	Ratio
5	−0.441	0.000	0.004	0.023	0.032
25	−0.207	0.001	0.011	0.042	0.084
50	−0.134	0.003	0.025	0.081	0.124
75	−0.070	0.006	0.044	0.146	0.206
95	−0.015	0.049	0.130	0.296	0.289
Mean	−0.167	0.008	0.039	0.106	0.143

we compute physical default probabilities, measures of variance risk premia, and Sharpe ratios. These quantities allow us to assess the economic soundness of our estimates beyond the in-sample fit. In Table IX, we provide summary statistics for a selection of economic quantities of interest. Our measure of the equity-level variance risk premium, *EVRP*, ranges between close to zero and −44%, averaging −17% per annum. The average *EVRP* is higher than our calibrated value of −11% for a representative Baa firm, although the median, at 13%, is closer.

Default probabilities average 0.8% for one year and 10.6% for 10 years. These numbers are somewhat higher than the historical average default rates reported by Moody's for Baa-rated firms during the 1920 to 2010 period of 0.29% and 6.91%, respectively. Note that our sample contains firms that are higher rated than Baa and firms that are lower rated than Baa. Historical default rates are asymmetric across ratings. In particular, historical default rates are almost three times higher for Ba than for Baa firms, while Baa rates are only twice as high as A rates. Taking an average of historical default rates using the rating distribution from our sample leads to a higher historical rate slightly greater than 8%. In addition, our distribution of estimated default probabilities is quite wide. Half of the one-year default probabilities lie between 0.1% and 0.6%. For the 10-year horizon, the range is 4.2% to 14.6%. The historical default rates are well within our estimated ranges for both horizons. Our median estimates of 0.3% and 8.1% are much closer to the historical rates. We conclude that our model's ability to fit spreads does not come at the expense of implying unreasonable physical default probabilities.

Figure 4 summarizes the model's empirical performance. The top two panels illustrate the time series of market and model spreads for 1- and 10-year maturities. In addition, they plot the model spreads with variance risk premia and jump components removed. Even with jumps, our model has trouble fitting one-year spreads on average both early in our sample, until 2004, and around

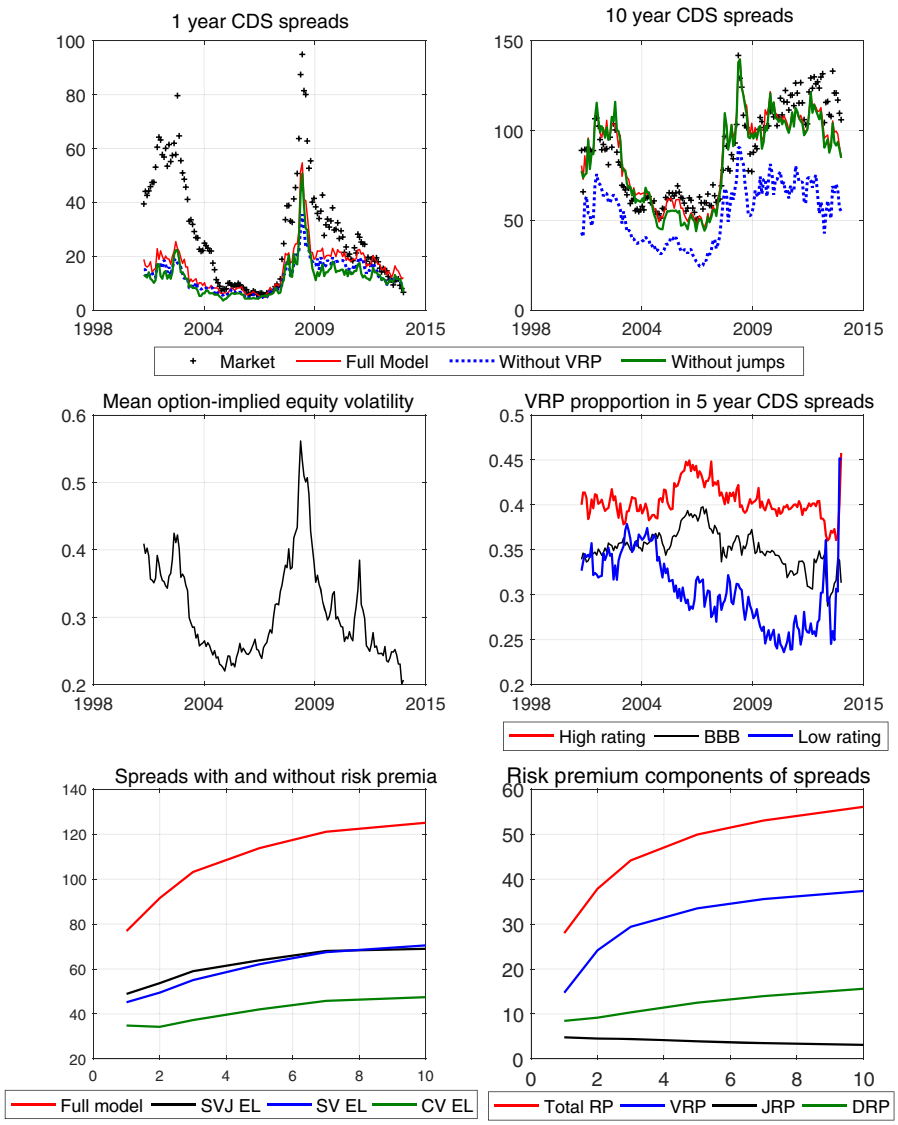

Figure 4. Credit spreads and risk premia. The two upper panels plot (for the 1- and 10-year tenors, respectively) mean market spreads, full model spreads, and spreads where we shut down the variance risk premium and jumps in the model. The middle left panel plots the average option-implied equity volatility. The middle right panel plots the proportion of variance risk premium (VRP) in spreads for firms rated A to AAA ("high rating"), those rated BBB, and those rated less than BBB ("low rating"). The lower left panel plots spreads with and without risk premia. The "Full" line represents the complete SVJ model with all risk premia as estimated. The lowest line represents the expected loss component of the term structure that would obtain in a constant volatility (CV) model, and the next two lines represent the expected loss components in the full model (SVJ) and the stochastic volatility model without jumps (SV). The bottom right panel plots the term structures of total risk premia, diffusive risk premia (DRP), variance risk premia (VRP), and jump risk premia (JRP). (Color figure can be viewed at wileyonlinelibrary.com)

the financial crisis. Shutting down jumps would have the most adverse effect on the level of model spreads, consistent with our calibration exercises. For 10-year default swaps, the model tracks levels of market spreads quite well, with the exception of a degree of underestimation after the financial crisis. Shutting down jumps would have a negligible effect on model spreads, while removing the variance risk premium from spreads would lead to a significant reduction in model spreads.

The middle left panel plots the average option-implied equity volatility. It is clear that the time series of average spreads peak at times of high volatility. The middle right panel illustrates the proportion of spreads due to variance risk premia over time for different credit ratings. Interestingly, in a relative sense, more highly rated firm spreads contain more of a variance risk premium—about 40% for firms rated A to AAA. BBB firm spreads contain about 35% variance risk premia, while the proportion for the speculative-grade firms in our sample is about another 5% lower.

The lower two panels consider average term structures of spreads and show how these depend on different types of risk premia. The lower left panel plots total spreads and expected loss spreads (spreads computed using P default probabilities) based on the CV, SV, and SVJ models. The expected loss spread when stochastic asset variance and jumps are removed is about 20 basis points lower than when they are present. The difference between the expected loss spreads with or without jumps is quite small but larger for the shorter maturities. On average, the risk premium makes up about half of the total spread. The lower right panel focuses on the risk premium components. Within the total risk premium, the largest component is the variance risk premium, the second is the diffusive premium that would obtain in a Black and Cox (1976) model, while the smallest component is the jump risk premium, which is the only component that is decreasing in maturity.

Overall, we find that the AVRP contributes significantly to empirical levels of spreads. Jump risk and the corresponding risk premium are smaller but important in explaining shorter maturity spreads.

VI. Extensions

A. Endogenous Default

In our baseline stochastic asset volatility model, we assume an exogenous fixed default boundary. This allows us to derive a quasi-closed-form solution. However, an exogenous default boundary ignores the influence of volatility on the optionality of equity and thus on the optimal default policy of a firm run in the interests of shareholders. We know from Leland (1994) that in the CV case, the endogenous default boundary is decreasing in the level of asset risk. The greater the volatility, the more valuable the optionality and hence shareholders will choose to default at a lower level of asset value. There is no reason that this intuition should not carry over to our setting. However, a complication that arises is that one of the key determinants of the boundary is stochastic.

To endogenize the optimal default policy, we solve the model by means of least squares Monte Carlo.[45] This allows us to compute, at any point in time and for any combination of asset value and variance, the value of servicing debt and keeping the firm alive, versus defaulting.[46] So instead of a default boundary being a line, it is a surface in the value and volatility dimensions. Figure 5 provides comparative statics based on the baseline parametrization in Table I. In the top panel, we examine the role of SV. The black line plots the boundary in Leland (1994) as a function of constant asset volatility. The other lines plot the default boundary for a given point in time as a function of realized asset volatility, given a long-run mean equal to the constant level in Leland (1994).

Consider the dashed line, which represents a combination of volatile asset volatility and slow mean reversion. The slope of the relationship has the same sign as in the Leland (1994) case but is lower, reflecting the fact that observing the level of asset volatility is now less informative than in the CV case. Increasing the speed of mean reversion works like a reduction in volatility and increases the boundary. It also slightly decreases the slope, suggesting that when mean reversion is strong, the particular level of the variance is less important than the long-run mean. Reducing the volatility parameter makes the boundary behave more similarly to that in the CV case.

Negative correlation between asset variance and value shocks decreases the level of the boundary, as on average volatility will be higher when approaching the boundary, thus increasing equity optionality. Finally, variance risk premia reduce the boundary. It is risk-adjusted volatility that matters for equity prices, so the presence of variance risk increases the patience of shareholders and reduces the boundary.

B. Capital Structure

We have shown that variance risk plays a role in the default decision. Variance risk benefits shareholders, all else equal. But the effect of a lower default boundary is dominated by the impact of increased risk-adjusted variance on the probability of hitting the boundary. In a trade-off model, this translates into higher risk-adjusted probabilities of default, which should increase distress costs and lower the value of the tax shield. Figure 6 illustrates this effect. It plots levered firm value as a function of leverage for two firms that are identical save the presence or absence of variance risk premia. Without variance risk premia, firm value reaches a peak at leverage a little higher than 70%. With variance risk premium, optimal leverage is reduced significantly, as is the

[45] See Longstaff and Schwartz (2001), who develop a tractable method to value American options by simulation. The main insight is that least squares regressions can be used to estimate the conditional expected value of continuation for a security holder faced with a decision to exercise an option or wait. The method is easy to apply in multifactor situations like the one at hand in this paper.

[46] McQuade (2018) also considers endogenous default in a similar setting. He finds that it helps him fit spreads in the cross section of credit ratings.

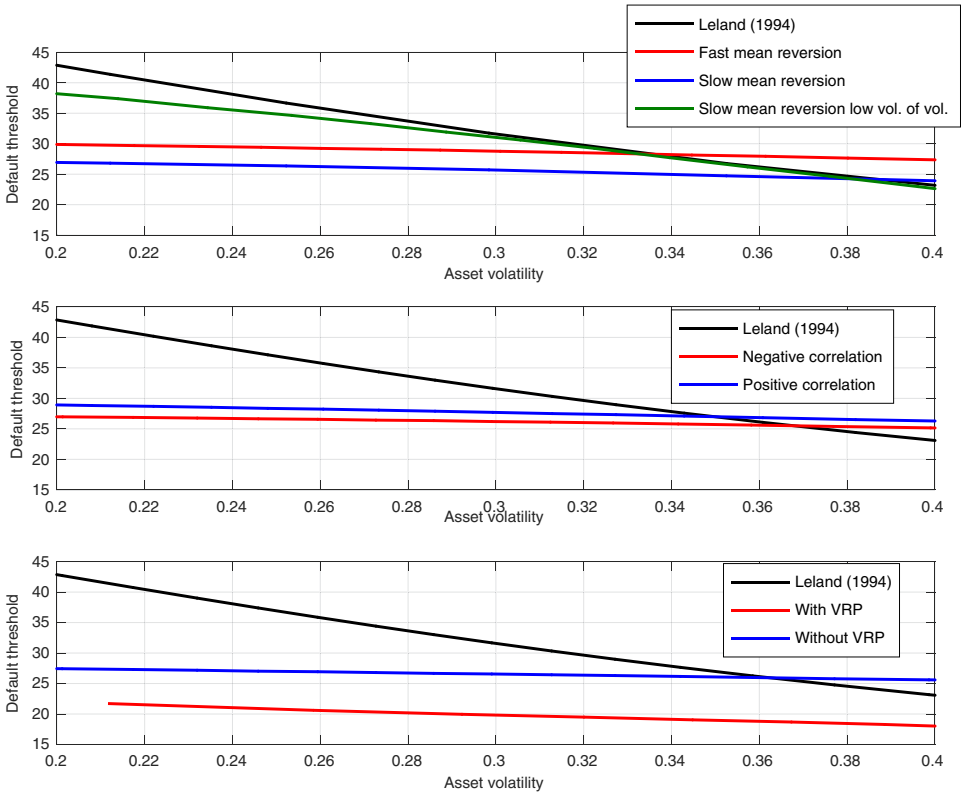

Figure 5. Stochastic volatilty, variance risk, and the firm's default policy. This figure illustrates how a firm's endogenous default policy depends on the firm's asset variance dynamics. The first two panels show how the default boundary depends on the parameters that govern the physical dynamics of asset variance. The final plot considers the influence of a variance risk premium (VRP) on the default decision. As a benchmark, we also plot the constant volatility boundary in Leland (1994), to which ours converges in the limit when asset variance becomes a constant. (Color figure can be viewed at wileyonlinelibrary.com)

level of levered firm value. These optimal leverage ratios are still high compared to the data. We assume a marginal tax rate of 35% and distress costs of 50% of asset value at default. These numbers are chosen to be consistent with Figure 7 in Leland (1994). More recent work uses lower tax rates (based on Graham (2000)) and lower bankruptcy costs (based on Andrade and Kaplan (1998), who report numbers about a third as large). It could be the case, however, that costs at default are in fact much higher ex ante than those observed ex post (see, e.g., Glover (2016)) or that distress costs are realized prior to default (see, e.g., Elkamhi, Ericsson, and Parsons (2012)). In addition, the dependence of bankruptcy costs on the state of the economy may be an alternative mechanism that influences the level of the optimal leverage. That said, while clearly only one among several potentially contributing factors, the presence of a

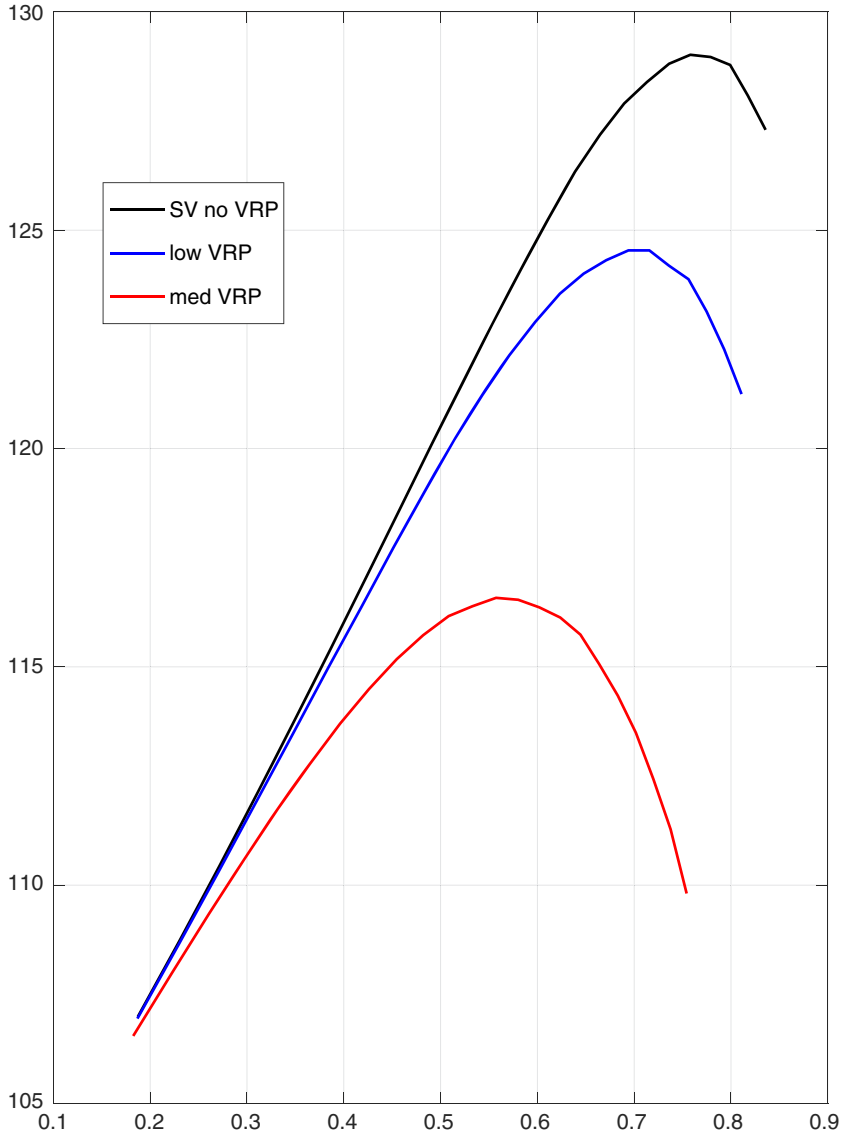

Figure 6. Stochastic volatility, variance risk, and the firm's leverage policy. The tax rate is assumed to be 35% and the proportional costs of financial distress, realized at default, amount to 50%. The initial asset variance is 0.2^2, the speed of mean reversion κ is 4, the correlation between the asset value and variance shocks ρ is -0.15, and the volatility of asset variance σ is 0.30. The default boundary is endogenous. (Color figure can be viewed at wileyonlinelibrary.com)

variance risk premium does have the potential to help resolve the low leverage puzzle.[47]

When we perform firm-by-firm estimation, we document significant heterogeneity in parameters related to jump and variance risk. Given we have shown that these risks influence firms' default and leverage decisions, it would be interesting to examine whether variation in these risks and the associated premia may help explain the empirical variation in firms' financing decisions.

VII. Conclusion

In this paper, we develop, evaluate, and estimate a first-passage-time structural credit risk model with priced SV to address the credit spread puzzle documented in Huang and Huang (2012) and further studied by CCG, among others. The key feature of this model is time variation in Sharpe ratios induced by a variance risk premium assumed to be proportional to asset variance. In addition, our model predicts lower Sharpe ratios for debt than equity, a finding suggestive of a link between the credit spread and distress premium puzzles.

We first evaluate our model by means of calibrations for representative Baa- and Aa-rated firms. We replicate the credit spread puzzle in the absence of variance risk and show that it can be resolved for medium- to longer-term maturities by allowing for a variance risk premium consistent with reasonable firm-level Sharpe ratios. Without jumps, however, the model is unable to fit shorter term credit spreads. Hence, in our empirical implementation we allow for jumps. Empirically, we study the ability of our model to jointly explain the dynamics of credit spreads and equity volatilities, a task that has been shown to be out of the reach of constant volatility structural credit risk models. The additional factor for asset variance permits our model to fit equity volatilities well, while significantly improving the fit for CDS prices relative to a nested constant volatility model. The estimation identifies economically significant variance risk premia and highlights that they explain an important part of spread levels.

We extend the model to study the role of variance risk in a firm's default and capital structure decisions. We find that shareholders pursue a more patient default policy (lower default boundary) but more conservative leverage policy in the presence of a variance risk premium.

The technical contribution of our paper—closed-form analytics for a first-passage-time SV model—has many obvious applications for the credit risk literature. More generally, we believe there are numerous applications in the real options literature, where investment and volatility are closely related.

Initial submission: January 17, 2013; Accepted: December 31, 2017
Editors: Bruno Biais, Michael R. Roberts, and Kenneth J. Singleton

[47] We repeated the exercise for various combinations of tax rates and distress costs. The results are qualitatively the same—optimal leverage ratios, and firm values at those leverage levels, are lower in the presence of a variance risk premium.

REFERENCES

Aït-Sahalia, Yacine, and Robert Kimmel, 2007, Maximum likelihood estimation of stochastic volatility models, *Journal of Financial Economics* 83, 413–452.

Andersen, Torben G., Tim Bollerslev, and Francis X. Diebold, 2007, Roughing it up: Including jump components in the measurement, modeling, and forecasting of return volatility, *Review of Economics and Statistics* 89, 701–720.

Andrade, Gregor, and Steven N. Kaplan, 1998, How costly is financial (not economic) distress? Evidence from highly levered transactions that became distressed, *Journal of Finance* 53, 1443–1493.

Bai, Jennie, and Pierre Collin-Dufresne, 2018, The determinants of the CDS-bond basis during the financial crisis of 2007-2009, *Financial Management*, forthcoming.

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

Bates, David, 2006, Maximum likelihood estimation of latent affine processes, *Review of Financial Studies* 19, 909–965.

Berndt, Antje, Rohan Douglas, Duffie Darrell, and Mark Ferguson, 2018, Corporate credit risk premia, *Review of Finance*, 22, 419–454.

Bhamra, Harjoat, Lars-Alexander Kuehn, and Ilya A. Strebulaev, 2010, The levered equity risk premium and credit spreads: A unified framework, *Review of Financial Studies* 23, 645–703.

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

Bongaerts, Dion, Frank de Jong, and Joost Driessen, 2011, Derivative pricing with liquidity risk: Theory and evidence from the credit default swap market, *Journal of Finance* 66, 203–240.

Brandt, Michael, and Pedro Santa-Clara, 2002, Simulated likelihood estimation of diffusions with an application to exchange rate dynamics in incomplete markets, *Journal of Financial Economics* 63, 161–210.

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

Campbell, John Y., Jens Hilscher, and Jan Szilagi, 2008, In search of distress risk, *Journal of Finance* 63, 2899–2939.

Cao, Charles, Fan Yu, and Zhaodong Zhong, 2010, The informational content of option-implied volatility for credit default swap valuation, *Journal of Financial Markets* 13, 321–343.

Chen, Hui, 2010, Macroeconomic conditions and the puzzles of credit spreads and capital structure, *Journal of Finance* 65, 2171–2212.

Chen, Long, Pierre Collin-Dufresne, and Robert S. Goldstein, 2009, On the relationship between the credit spread puzzle and the equity premium puzzle, *Review of Financial Studies* 22, 3367–3409.

Cheredito, Patrick, Damir Filipović, and Robert L. Kimmel, 2007, Market price of risk specifications for affine models: Theory and evidence, *Journal of Financial Economics* 83, 123–170.

Choi, Jaewon, and Matthew Richardson, 2016, The volatility of a firm's assets and the leverage effect, *Journal of Financial Economics* 121, 254–277.

Collin-Dufresne, Pierre, and Robert S. Goldstein, 2001, Do credit spreads reflect stationary leverage ratios? *Journal of Finance* 56, 1929–1957.

Collin-Dufresne, Pierre, Robert S. Goldstein, and Fan Yang, 2012, On the relative pricing of long-maturity index options and collateralized debt obligations, *Journal of Finance* 67, 1983–2014.

Cox, John C., Jonathan E. Ingersoll, and Stephen A. Ross, 1985, A theory of the term structure of interest rates, *Econometrica* 53, 385–407.

Dai, Qiang, and Kenneth J. Singleton, 2000, Specification analysis of affine term structure models, *Journal of Finance* 55, 1943–1978.

Davydenko, Sergei A., 2012, When do firms default? A study of the default boundary, Working paper, University of Toronto.

Dennis, Patrick, Stewart Mayhew, and Chris Stivers, 2006, Stock returns, implied volatility innovations, and the asymmetric volatility phenomenon, *Journal of Financial and Quantitative Analysis* 41, 381–406.

Doshi, Hitesh, Jan Ericsson, Kris Jacobs, and Stuart Turnbull, 2013, Pricing credit default swaps with observable covariates, *Review of Financial Studies* 26, 2049–2094.
Driessen, Joost, 2005, Is default event risk priced in corporate bonds? *Review of Financial Studies* 18, 165–195.
Duffee, Gregory R., 2002, Term premia and interest rate forecasts in affine models, *Journal of Finance* 57, 405–443.
Duffie, Darrell, and David Lando, 2001, Term structures of credit spreads with incomplete accounting information, *Econometrica* 69, 633–664.
Duffie, Darrell, and Kenneth J. Singleton, 1999, Modeling term structures of defaultable bonds, *Review of Financial Studies* 12, 687–720.
Elkamhi, Redouane, Jan Ericsson, and Christopher A. Parsons, 2012, The cost and timing of financial distress, *Journal of Financial Economics* 105, 62–81.
Elkamhi, Redouane, and Chayawat Ornthanalai, 2010, Market jump risk and the price structure of individual equity options, Working paper, University of Toronto.
Eom, Young Ho, Jean Helwege, and Jing-Zhi Huang, 2004, Structural models of corporate bond pricing: An empirical analysis, *Review of Financial Studies* 17, 499–544.
Eraker, Bjørn, Michael S. Johannes, and Nick Polson, 2003, The impact of jumps in volatility and returns, *Journal of Finance* 58, 1269–1300.
Feldhütter, Peter, and Stephen M. Schaefer, 2018, The myth of the credit spread puzzle, *Review of Financial Studies* 31, 2897–2942.
Fortet, Robert, 1943, Les fonctions aléatoires du type de Markoff associées à certaines équations linéaires aux dérivées partielles du type parabolique, *Journal de Mathématiques Pures et Appliquées* 22, 177–243.
Fouque, Jean-Pierre, George C. Papanicolaou, and Ronnie Sircar, 2000, *Derivatives in Financial Markets with Stochastic Volatility* (Cambridge University Press, Cambridge, New York).
Fouque, Jean-Pierre, George C. Papanicolaou, Ronnie Sircar, and Knut Sølna, 2003, Singular perturbations in option pricing, *SIAM Journal on Applied Mathematics* 63, 1648–1665.
Fouque, Jean-Pierre, George C. Papanicolaou, Ronnie Sircar, and Knut Sølna, 2004, Stochastic volatility corrections for interest rate derivatives, *Mathematical Finance* 14, 173–200.
Fouque, Jean-Pierre, George C. Papanicolaou, Ronnie Sircar, and Knut Sølna, 2006, Stochastic volatility effects on defaultable bonds, *Applied Mathematical Finance* 13, 215–244.
Fouque, Jean-Pierre, George C. Papanicolaou, Ronnie Sircar, and Knut Sølna, 2011, *Multiscale Stochastic Volatility for Equity, Interest Rate, and Credit Derivatives* (Cambridge University Press, Cambridge, New York).
Garlappi, Lorenzo, Tao Shu, and Hong Yan Yan, 2008, Default risk, shareholder advantage and stock returns, *Review of Financial Studies* 21, 2743–2778.
Garlappi, Lorenzo, and Hong Yan, 2010, Financial distress and the cross-section of equity returns, *Journal of Finance* 66, 789–822.
Glover, Brent, 2016, The expected cost of default, *Journal of Financial Economics* 119, 284–299.
Graham, John, 2000, How big are the tax benefits of debt? *Journal of Finance* 55, 1901–1942.
Han, Bing, and Yi Zhou, 2011, Variance risk premium and cross-section of stock returns, Working paper, UT Austin.
Heston, Steven L., 1993, Closed-form solution for options with stochastic volatility, with application to bond and currency options, *Review of Financial Studies* 6, 327–343.
Huang, Jing-Zhi, and Ming Huang, 2012, How much of the corporate-treasury yield spread is due to credit risk? *Review of Asset Pricing Studies* 2, 153–202.
Huang, Jing-Zhi, and Hao Zhou, 2008, Specification analysis of structural credit risk models, Working paper, Pennsylvania State University.
Huang, Xin, and Gorge Tauchen, 2005, The relative contribution of jumps to total price variance, *Journal of Financial Econometrics* 3, 456–499.
Jones, E. Philip, Scott P. Mason, and Eric Rosenfeld, 1984, Contingent claims analysis of corporate capital structures: An empirical investigation, *Journal of Finance* 39, 611–625.
Leland, Hayne E., 1994, Corporate debt values, bond covenants, and optimal capital structure, *Journal of Finance* 49, 1213–1252.

Leland, Hayne E., 2006, *Princeton Lectures in Finance* (Princeton University Press, Princeton, NJ).

Longstaff, Francis A., and Edouardo S. Schwartz, 1995, A simple approach to valuing risky fixed and floating rate debt, *Journal of Finance* 50, 789–821.

Longstaff, Francis A., and Edouardo S. Schwartz, 2001, Valuing American options by simulation: A simple least-squares approach, *Review of Financial Studies* 14, 113–147.

Lustig, Hanno, and Adrien Verdelhan, 2012, Business cycle variation in the risk-return trade-off, *Journal of Monetary Economics* 59, 35–49.

Marra, Miriam, 2017, Explaining co-movements between equity and CDS bid-ask spreads, *Review of Quantitative Finance and Accounting* 49, 811–853.

McQuade, Timothy, J., 2018, Stochastic volatility and asset pricing puzzles, Working paper, Stanford University.

Merton, Robert C., 1974, On the pricing of corporate debt: The risk structure of interest rates, *Journal of Finance* 29, 449–470.

Miranda, Mario J., and Paul L. Fackler, 2002, *Applied Computational Economics and Finance* (MIT Press, Cambridge, MA).

Moody's Investor Service, 2010, *Corporate Default and Recovery Rates* (Moody's Corporation, New York). Available at https://www.moodys.com/.

Pan, Jun, 2002, The jump-risk premia implicit in options: Evidence from an integrated time-series study, *Journal of Financial Economics* 62, 3–50.

Pan, Jun, and Kenneth J. Singleton, 2008, Default and recovery implicit in the term structure of sovereign CDS spreads, *Journal of Finance* 63, 2345–2384.

Wang, Hao, Hao Zhou, and Yi Zhou, 2013, Credit default swap spreads and variance risk premia, *Journal of Banking and Finance* 37, 3733–3746.

Zhang, Benjamin Y., Hao Zhou, and Haibin Zhu, 2009, Explaining credit default swap spreads with equity volatility and jump risks of individual firms, *Review of Financial Studies* 22, 5099–5131.

Supporting Information

Additional Supporting Information may be found in the online version of this article at the publisher's website:

Appendix S1: Internet Appendix.