Implied Default Probabilities and Losses Given Default from Option Prices*

Jennifer Conrad[†]
Robert F. Dittmar[‡]
Allaudeen Hameed[§]

February 15, 2020

Abstract

We propose a novel method of estimating default probabilities using equity option data. The resulting default probabilities are highly correlated with estimates of default probabilities extracted from CDS spreads, which assume constant losses given default. Additionally, the option-implied default probabilities are higher in bad economic times and for firms with poorer credit ratings and financial positions. A simple inferred measure of loss given default is related to underlying business conditions, and varies across sectors; the time-series properties of this measure is similar after controlling for liquidity effects.

*This paper has benefitted from the comments of Xudong An, Nina Boyarchenko, Davie Henn, Ken Singleton, Yin-Hua Yeh, and Adam Zawadowski, as well as seminar participants at the 2010 Financial Economics in Rio conference at FGV Rio de Janeiro, the 2011 FMA Asian conference, the 2012 American Finance Association conference, the 2012 European Finance Association conference, the 2013 NUS-RMI conference, the 2017 China International Review conference, Aalto, Arizona State, Georgetown, Goethe, Indiana, Purdue, and Rice Universities, Hong Kong University of Science and Technology, the Stockholm School of Economics, and the Universities of British Columbia, Mannheim, Toronto, Utah, and Western Ontario. Hameed gratefully acknowledges financial support from NUS Academic Research Grant. All errors are the responsibility of the authors.

[†]Department of Finance, Kenan-Flagler Business School, University of North Carolina, j_conrad@kenan-flagler.unc.edu

[‡]Department of Finance, Ross School of Business, University of Michigan, rdittmar@umich.edu

[§]Department of Finance, NUS Business School, National University of Singapore, allaudeen@nus.edu.sg

1 Introduction

From the perspective of academics, market professionals, and regulators, one of the attractive features of a credit default swap (CDS) contract is its window into market perceptions of credit risk. Based on no-arbitrage pricing formulations and assumptions about the recovery rate on the asset, one can use information in the quoted spread on a CDS contract to infer the market's implied risk neutral probability of default; see, e.g., Duffie and Singleton (1999). This measure stands as a market-based nonparametric alternative to agency credit ratings and structural models of default. However, the market for CDS reached its peak in late 2007; since that time, this market has substantially decreased in size.[1] Recent market events suggest that the reduction in the size of the CDS market will not be reversed soon, as noted in Augustin, Subrahmanyam, Tang, and Wang (2016). Thus, the decline in the single-name CDS market represents a loss for those interested in market-driven estimates regarding the credit risk of a corporate entity.

In this paper, we propose a novel measure of the risk neutral probability of default based on option prices. It is well-known that option prices are informative about the risk neutral distribution of equity payoffs; see, e.g., Breeden and Litzenberger (1978). In addition, the equity payoff depends on default risk. In principle, if absolute priority holds, the value of equity will be zero in the case of a default as in Merton (1974). Alternatively, one can define a default region of the equity payoff distribution and use the cumulative probability of that region inferred from option prices to identify a probability of default. Importantly, an option-based approach has the advantage that, since options are traded on a large number of underlying securities, this method could provide estimates of default probability for a broader cross-section of firms.

Our results indicate that estimates of the levels of implied default probabilities extracted from equity options are strongly, but not perfectly, correlated with default probabilities estimated using CDS. If we assume constant loss given default (as is frequently done in the CDS market), the median correlation between estimates of default probabilities for the cross-section of firms extracted from

[1]The Bank for International Settlements reported notional principal outstanding in the last half of 2007 of \$61 trillion; in the first half of 2019, this figure stood at \$7.8 trillion.

these two markets is 0.52. Aggregated across firms, the two estimates of default probabilities are highly correlated through time, with a correlation coefficient of 0.66. The default probabilities estimated from equity options prices increase monotonically with lower credit ratings; in addition, the relations between default probabilities estimated from equity options and firm characteristics are similar to the relations estimated between CDS default probabilities and firm characteristics. Overall, the evidence suggests that equity options can provide important information concerning the probability of default for a broad spectrum of underlying firms.

We also investigate whether the imperfect correlation between option- and CDS-implied default probabilities reflects the fact that estimates of loss given default (LGD) embedded in CDS rates vary across firms and through time, as argued in Berndt, Douglas, Duffie, and Ferguson (2018), Doshi, Elkamhi, and Ornthanalai (2018), Schuermann (2004) and Altman, Bradi, Resti, and Sironi (2005). We find that a simple proxy for LGD, defined as the ratio of the CDS spread to option-implied default probability, covaries positively with the frequency of default in the aggregate sample; the ratio is high during the early 2000s and the financial crisis, declines in the mid-2000s and after the financial crisis and has been relatively low during 2013-2017. While some of this difference in the option-implied default probability and CDS spread is related to measures of illiquidity, the time-series patterns in this variable remain after removing variation due to illiquidity.

Our paper is related to the literature investigating the ability of option prices to provide information about default. In an earlier paper whose intuition is closely related to our work, Le (2015) develops models of option and CDS prices, and uses those models sequentially to estimate default probabilities and recovery rates. His method of estimating default probabilities differs from ours; for example, he assumes a specific process for the dynamics of equity prices, and estimates the probability of default, where default is associated with equity prices diffusing or jumping to zero. In contrast, we estimate the risk-neutral density at a particular point in time from a range of option prices and consider various default thresholds. Moreover, our sample (which extends through February 2017) enables us to document the time series properties of loss given default over a longer period, including the credit crisis. Capuano (2008) uses a cross-entropy functional to infer default probabilities, although Vilsmeier (2014) notes that the entropy approach has issues with accuracy and numerical stability (and provides some technical fixes for those problems); our approach is simpler computationally. Carr and Wu (2011) show that deep out-of-the-money options can be used to synthesize a default insurance contract, and, as a consequence, infer the probability of default. Their approach is simple and intuitive, but necessitates the existence of options that are very deep out of the money, limiting the number of firms for which these probabilities can be calculated.

This paper is also related to others that have inferred losses given default (LGD) or other credit-related measures from the market prices of various securities; the papers in this literature differ with regard to the securities that are required for estimation of LGD, and, in some cases, the processes that rates of LGD are assumed to follow. Duffie and Singleton (1999) and Das and Sundaram (2007) provide examples of how LGD might be inferred from securities with the same probability of default but different payout structure or priority. Similarly, Madan and Unal (1998) empirically investigate the separation of default probability and LGD using junior and senior debt prices where those are available; Madan, Unal, and Güntay (2003) also exploit differences in debt priority to infer losses given default. Doshi, Elkamhi, and Ornthanalai (2018), as noted above, use the term structure of CDS to estimate LGD. Bakshi, Madan, and Zhang (2006) use a risky debt model with stochastic rates of LGD to infer measures of LGD from risky bond prices. More recently, Berndt, Douglas, Duffie, and Ferguson (2018) combine information from CDS rates, firm-specific default estimates from Markit, and information regarding expected default from Moody's Analytics to estimate credit risk premia.

The remainder of the paper is organized as follows. In Section 2, we discuss the methodology we employ for extracting risk neutral default probabilities from options and from CDS spreads. We describe the data that we employ in this paper in section 3, and present estimates of default probabilities and their relation to various firm characteristics. In Section 4, we explore the differences between the two measures of default probability, including cross-sectional and time-series variation in LGD and liquidity effects. We conclude in Section 5.

2 Risk Neutral Probabilities Implied by CDS and Option Prices

2.1 Pricing Credit Default Swaps

In order to infer risk neutral default probabilities from the prices of credit default swaps (CDS), we follow a model widely used in practice for their valuation, detailed in O'Kane and Turnbull (2003). In the discussion that follows, we assume that the swap being valued is a one-year CDS contract with quarterly premium payments, and that there is no information on CDS from which to infer risk neutral default probabilities for horizons of less than one year. Under these assumptions, the practice is to assume a flat default probability term structure over the year.[2]

When a CDS contract is struck, the swap premium is set such that the value of the premium leg, received by the writer of the swap, is equal to the value of the protection leg, received by the swap purchaser. Assuming that premiums are accrued in case of default during a quarter, the value of the premium leg is given by

$$\frac{1}{2}s_t \sum_{j=1}^{4} 0.25 e^{-r\left(\frac{j}{4}\right)\frac{j}{4}} \left(e^{-\lambda_t \frac{j-1}{4}} + e^{-\lambda_t \frac{j}{4}}\right), \tag{1}$$

where $r(\tau)$ is the continuously compounded zero coupon Treasury yield with maturity τ expressed as an annual rate and λ_t is the default intensity. Because of the assumption of a flat term structure of default probability over one year, this intensity is invariant to maturity for the one-year horizon, although it is indexed by t to indicate that default intensity may change over time. The quantity $e^{-\lambda_t \tau}$ represents the risk neutral probability that the entity survives to time τ. Intuitively, expression (1) simply calculates the present value of the swap payments received by the swap writer, conditional on survival of the entity.

[2]Note that this valuation model is similar to the one described in Augustin and Saleh (2017), with the exception that we adjust for accrued interest.

The value of the protection leg is the risk neutral expected loss on the CDS,

$$(LGD)\sum_{j=1}^{12} e^{-r(\frac{j}{12})\frac{j}{12}}\left(e^{-\lambda_t\frac{j-1}{12}} - e^{-\lambda_t\frac{j}{12}}\right), \tag{2}$$

where LGD designates the loss given default as a fraction of the amount owed. Expression (2) is a discrete approximation to an integral that represents the expected risk neutral loss on the underlying entity. O'Kane and Turnbull (2003) show that for a constant default intensity, the approximation error is given by $\frac{r(\tau)}{2M}$, where $r(\tau)$ is the continuously compounded risk free rate over the constant default intensity horizon and M is the number of summation periods. The authors suggest that for $M = 12$ as above and a risk free rate of 3%, the absolute value of the error is 1 basis point on a spread of 800 basis points.

The above expressions allow for the determination of the break-even CDS spread if one has estimates of losses given default and risk neutral survival probabilities. Alternatively, the expressions can be used to infer risk neutral default probabilities given rates of LGD and constant maturity CDS spreads. For example, using the expressions above, the break-even CDS spread is given by

$$s_t = \frac{(LGD)\sum_{j=1}^{12} e^{-r(\frac{j}{12})\frac{j}{12}}\left(e^{-\lambda_t\frac{j-1}{12}} - e^{-\lambda_t\frac{j}{12}}\right)}{\frac{1}{2}\sum_{j=1}^{4} 0.25e^{-r(\frac{j}{4})\frac{j}{4}}\left(e^{-\lambda_t\frac{j-1}{4}} + e^{-\lambda_t\frac{j}{4}}\right)}. \tag{3}$$

In the absence of any market frictions, equation(3) is a nonlinear equation in the default intensity, λ_t.

For our initial calculations, we assume a constant rate of $LGD = 0.60$, consistent with common practice. Given data on the risk-free term structure and one-year credit default swaps, we then solve equation (3) for λ_t for each reference entity and date in our sample. Given this default intensity, the probability of default of entity i is given by

$$Q_{i,t}^C = 1 - e^{-\lambda_{i,t}}, \tag{4}$$

where the superscript C indicates that CDS data are used to infer default probabilities.

2.2 Measuring Default Probabilities from Option Prices

An options-based approach to extracting default probabilities is based on the seminal work of Breeden and Litzenberger (1978), who show that one can recover the risk neutral density of equity returns from option prices. Given the risk neutral density, the risk neutral probability of default can be thought of as the mass under the density up to the return that corresponds to a default event.

We construct the risk-neutral density using estimates of the risk neutral moments computed as in Bakshi, Kapadia, and Madan (2003) and the Normal Inverse Gaussian (NIG) method developed in Eriksson, Ghysels, and Wang (2009). Specifically, Bakshi, Kapadia, and Madan (2003) (BKM) show that one can use traded option prices to compute estimates of the variance, skewness, and kurtosis of the risk neutral distribution. These moments in turn serve as the inputs to the NIG method, which estimates the distribution. Eriksson, et al show that the NIG has several advantages to alternatives such as Gram-Charlier series expansions in pricing options. In particular, the distribution prevents negative probabilities, which the expansions can generate for the levels of skewness and kurtosis implied by option prices. The density is also known in closed form, avoiding the computational intensity of expansion approaches. Details of the estimation process are provided in Online Appendix A.

Once the risk-neutral distribution is estimated, we measure the probability of default for entity i at time t as the cumulative density of the NIG distribution at a default threshold α,

$$Q^O_{it}(\tau) = \int_{-\infty}^{\alpha} f_{NIG}\left(x, \mathcal{E}_{i,t}(\tau), \mathcal{V}_{it}(\tau), \mathcal{S}_{it}(\tau), \mathcal{K}_{it}(\tau)\right) dx \tag{5}$$

where f_{NIG} is the NIG density function evaluated at a log return of x with parameters calculated as in BKM.[3] The superscript O in equation (5) indicates that the risk neutral probability has been recovered from options data. The exact functional form of the density is provided in the Online

[3] In a few cases, our estimates of the kurtosis are too small given the calculated skewness. In order to calculate the cumulative density, it is necessary that $K_{it} > 3 + \frac{5}{3}S_{it}^2$. In cases in which this restriction is violated, we set the kurtosis to $K_{it} = 3 + \frac{5}{3}S_{it}^2 + 1e-14$.

Appendix A.

A critical detail in this procedure is the definition of the default threshold, α. In the Merton (1974) model, equity has zero value in the case of default. However, the density at $x = \ln(0)$ cannot be calculated. Carr and Wu (2011) deal with this problem by assuming that there is a range of values for the stock price, $[A, B]$, in which default occurs. Specifically, prior to default, the equity value is assumed to be greater than B, and upon default the value is assumed to drop below the value $A \in [0, B)$. In their empirical implementation, the authors set $A = 0$, and choose the lowest priced put with positive bid price and positive open interest with strike price less than \$5 and option delta of less than or equal to 15% in absolute value for their estimation.

For our initial choice of the default threshold, we begin with an updated version of the data on bankruptcy filing dates from Chava and Jarrow (2004).[4] We merge these data with CRSP, and calculate the percentage decline in price ΔP from 12 months prior to the bankruptcy filing to either the delisting date or the CRSP price observed in the bankruptcy month.[5]

We examine the extent to which these declines are related to other measures of credit risk and to credit ratings. In this sample, there are 383 firms for which we have S&P long-term credit ratings for the borrowers 12 months prior to bankruptcy. The 12-month average declines in price for these firms are depicted by credit rating in Table 1, where we group together all firms of a particular letter grade (i.e., 'A', 'A+', and 'A-'). The table suggests that there is a clear relation between credit rating and price decline in the 12 months leading up to bankruptcy. Between 'BBB'-rated and 'CC'-rated firms, there is a near monotonic decline in losses. The prices of 'BBB'-rated firms are on average 8% of their 12-month prior levels and the prices of 'CC'-rated firms are 45% of their 12-month prior levels on average. The magnitude of price declines does not increase perfectly with credit rating; the average price decline of 'A'-rated firms is higher than that of 'BBB'-rated firms and the price decline of 'BB'-rated firms is higher than that of 'B'-rated firms.[6] However,

[4] Thanks to Sudheer Chava and Claus Schmitt for making these data available.

[5] There are 1560 bankruptcy events in which the price declined over the previous twelve months, with an average decline in price of 79.40%. There are an additional 86 cases in which returns are positive over the 12 month period.

[6] The average price declines of 'A'-rated firms are driven by the very small number of 'A'-rated firms (3) that have filed for bankruptcy. These firms are PG&E, with a stock price drop of 63.6% prior to filing in April, 2001; Armstrong Cork, with a stock price drop of 93.8% prior to filing in December, 2000; and Lehman Brothers, with a

the overall pattern suggests that average price declines in the 12 months leading up to bankruptcy filing decrease with credit rating.

Based on this evidence, for much of the analysis in the paper we let α, the default threshold, vary across credit ratings; we term this the Rating-derived α. According to Standard and Poor's, from 1981-2015, no AAA-rated credit has defaulted, and defaults of 'AA'-rated are rare.[7] As a consequence, we assign $\alpha = 0.05$ to 'AAA'-rated firms and $\alpha = 0.10$ to 'AA'-rated firms respectively. We assign $\alpha = 0.15$ to 'A'-rated firms and increment α by 0.05 as we move to lower levels of credit ratings from 'BBB' to 'B', consistent with the lower price declines associated with these firms. Finally, firms with ratings of 'CCC' and below are assigned $\alpha = 0.35$. In our sample, these choices imply an average default threshold of $\alpha = 0.18$.

In robustness checks, we also compute default probabilities for each firm assuming a constant critical value across firms, and allow those critical values to vary from $\alpha = 0.01$ to $\alpha = 0.40$. In later sections of the paper, we compare results obtained using Rating-derived alphas to those obtained using a constant threshold of 0.15, and all results with constant thresholds are available from the authors upon request.

3 Risk Neutral Probabilities Implied by CDS and Option Prices

3.1 Data Description

Data on CDS are obtained from Markit. The initial sample consists of daily representative CDS quotes on all entities covered by Markit over the period January 2002 through November 2019. With the standardization of CDS contracts in 2009, new CDS contracts began trading with fixed coupon of 100 or 500 basis points, with upfront payment depending on the perceived credit risk of the underlying bond issuer. The CDS rate provided by Markit are 'at market' composite CDS quotes, computed based on the bid and ask quotes obtained from two or more anonymous CDS

drop of 99.9% prior to delisting in September, 2008.

[7] These data are sourced from Standard and Poor's Annual Global Corporate Default Study and Ratings Transitions 2015.

dealers. We assume that the composite CDS rate is the rate at which the market value of the default swap is zero, without an upfront payment. While the five-year contract is generally thought to be the most liquid, our proposed measure of default probability relies on options data, of which few are struck for maturities in excess of one year. As a consequence, we restrict attention to entities which have quote data available on one year CDS. We use these quoted prices, together with zero coupon discount rates, to solve for the default intensity, λ_t in equation (3), assuming a constant rate of LGD at 60%. Discount rates are obtained by fitting the extended Nelson and Siegel (1987) model in Svensson (1994) using all non-callable Treasury securities from CRSP. Our initial sample consists of 364 entities for which we have at least one default intensity observation.

Options data are from OptionMetrics. The calculation of the risk neutral moments requires the computation of integrals over a continuum of strikes. However, options are struck at discrete intervals. In addition, while the CDS in our sample have a constant one-year maturity, the maturity of options available in our sample varies and there are relatively few contracts available that are close to one year to maturity. We follow Hansis, Schlag, and Vilkov (2010) and Chang, Christoffersen, and Jacobs (2013) in constructing the volatility surface for options at 365 days to maturity using a cubic spline. We interpolate implied volatilities over the support of option deltas ranging from -99 to 99 at one-delta intervals, setting implied volatilities constant for deltas outside the span of observed option prices. We then convert implied volatilities to option prices and integrate over out-of-the-money calls and puts using the rectangular approximation in Dennis and Mayhew (2002). In order to be included in the sample, we require that options have positive open interest, positive bid and offer prices, at least two out-of-the-money puts and out-of-the-money calls, offer prices greater than bid prices and offer prices greater than $0.05. We also eliminate options where the offer price is greater than five times the bid price. Our sample of option-implied default probabilities yields 299 firms out of the 364 firms with default intensity information from CDS data.

We merge data from Markit, OptionMetrics and CRSP data on the basis of the ticker and firm name to obtain the permno as a unique identifier for each firm. Next, we merge the matched sample of option- and CDS-implied default probabilities with credit ratings data from Compustat.

We retain observations for which Compustat has a Standard and Poor's ratings grade for the month of the observation. Since credit ratings data on Compustat end in February 2017, the combined sample extends only through that date. The final sample of firms, which have at least one time series observation with an option-implied default probability, a CDS-implied default probability, and a Standard and Poor's credit rating, consists of 276 firms over the period January 2002 through February 2017.

3.2 Descriptive Statistics

Summary statistics for the default probabilities implied by options and credit default swaps in our sample of firms are presented in Table 2. For each firm, the default probability is calculated as the time-series average of the weekly estimates obtained from CDS or options data using Rating-derived alphas. We report the mean, standard deviation, fifth, fiftieth and ninety-fifth percentiles of the distribution of default probabilities in Panel A. Additionally, we report the fifth, fiftieth, and ninety-fifth percentile of the distribution of the correlation between CDS- and option-implied probabilities in Panel B.

The summary statistics indicate that across the distribution of firms, option-implied probabilities are on average higher than CDS-implied default probabilities. The mean and median option-implied default probability are 2.87% and 2.03%, compared to 1.98% and 0.77% for CDS. In addition, the option-implied probabilities exhibit less cross-sectional variation than CDS-implied default probabilities, with standard deviations of 2.61% and 3.75%, respectively. These results suggest that the distribution of CDS-implied default probabilities is more positively skewed than option-implied default probabilities; we explore various explanations for the skew later in the paper.

We observe a substantial but imperfect correlation between CDS-implied and option-implied probabilities. The median correlation coefficient between default probabilities is 0.52; even at the 95th percentile, the correlation is less than 1.0.[8]

[8] At the fifth percentile, note that the correlation between the two default probabilities is negative. However, we find that this result is driven by the large gaps in the time series for some firms, with virtually all of the negative correlations concentrated in firms with relatively few time series observations.

In Figure 1, we plot the time series of cross-sectionally averaged weekly default probabilities implied by CDS and options. For the options-based default probabilities, we show the averages when the default threshold varies by credit rating, as well as when the default threshold is held constant at 0.15. The plot shows that the CDS measure of default probability, as well as both option-implied default probabilities, exhibit common features; default probabilities are uniformly low during the economic expansion and spike during times of economic turbulence. In particular, the default probabilities rise sharply during the recession in the early 2000s and the financial crisis of 2007-2009. Default probabilities also spike in late 2011, corresponding to the uncertainty surrounding the United States Congress' willingness to raise the federal debt ceiling and the subsequent downgrade of U.S. sovereign debt by Standard and Poor's. Both measures of default probabilities are low during the economic expansions of mid-2000s and the later part of the sample in 2013-2017. The result that default probabilities are countercyclical is consistent with the evidence in Chen, Collins-Dufresne and Goldstein (2008), who find that they are better able to explain Baa-Aaa spreads using a model that is calibrated to match default rates that are countercyclical. We find that all three estimates of default probabilities exhibit this tendency: the correlation between both of the two option-implied default probabilities and the CDS implied default probability through time is relatively high, at 0.66 and 0.56, respectively.

The greater variability in CDS-implied default probabilities observed in Figure 1, with higher default probabilities compared to option-implied probabilities during market downturns and lower default probabilities compared to option-implied probabilities during market upturns, is consistent with the evidence in Table 2. It may also reflect variation in losses given default through the cycle, in contrast with the (assumed) constant losses given default embedded in the CDS-implied estimates of default probabilities shown here. That is, this pattern is consistent with expected LGD covarying positively with true default probabilities. We investigate this possibility later in the paper.

3.3 Default Probabilities and Credit Ratings

To investigate further the cross-sectional variation in implied default probabilities, we report summary statistics for default probabilities by credit rating. In Table 3, we present mean and median default probabilities and the number of firms conditional on ratings class. Since firms may migrate across ratings, the total N reported in the table differs from the total number of firms reported in the sample. Thus, the $N = 9$ for 'AAA'-rated firms indicates that there are 9 firms that at some point in this time series have been rated 'AAA.' As above, we group together firms with a '+,' '-,' or no modifier.

The summary statistics in Table 3 indicate that for both CDS- and option-implied default probabilities, the average risk neutral probability of default increases monotonically across ratings classes. Using options (CDS) data, average default probabilities increase from 0.45% (0.23%) for 'AAA'-rated firms to 14.39% (15.38%) for firms rated 'CCC+' and below. There is a similar pattern of monotonic increase in median default probabilities across ratings classes, although CDS-implied median default probabilities are typically lower than the mean default probabilities. Across all ratings except 'CCC', the average and median option-implied probability is higher than that of the CDS-implied probability, consistent with the aggregate evidence reported earlier.

Of course, since the default thresholds in Table 3 vary by credit rating, some of these results may be mechanically induced. In untabulated results, we examine default probabilities across ratings classes, while keeping the default threshold constant across firms. Our results indicate that, as credit ratings deteriorate, the default probabilities implied by options for firms with poor credit ratings and relatively high default thresholds (high α) are more similar to those implied by CDS than those implied by relatively low thresholds (low α). Similarly, the default probabilities implied by options for firms with strong credit ratings are more similar to those implied by CDS when the threshold is low. This evidence is consistent with default thresholds that vary both cross-sectionally (as in Chava and Jarrow (2004)) and over time (as in Chen, Collins-Dufresne and Goldstein (2008)) as credit ratings migrate.

3.4 Firm Characteristics and Probability of Default

The evidence presented above indicates that option-implied risk neutral probabilities of default are substantially and positively correlated with CDS-implied risk neutral probabilities of default, and are cross-sectionally correlated with Standard and Poor's credit ratings. We analyze cross-sectional variation in default probabilities in more detail in this section, by examining the relation between both estimates of default probability and firm characteristics. In particular, we use a variant of Campbell, Hilscher, and Szilagyi (2008), who specify a pooled logit model for prediction of default. Instead of using a limited dependent variable based on an observation of default, we regress continuous estimates of default probabilities on firm-specific variables.

$$Q_{it}^{k} = a_{it} + \mathbf{b}_{it}'\mathbf{x}_{it} + u_{it},$$

where Q_{it}^{k} is the week t observation of the risk neutral probability, $k = \{C, O\}$ indexes CDS- and option-implied default probabilities and $\mathbf{x}_{it}$ is a vector of firm-specific characteristics. In constructing these characteristics, we sample market variables at the weekly frequency contemporaneously with the default probabilities. Accounting data items are sampled at the quarterly frequency and lagged one quarter relative to the default probabilities.[9] The firm-specific variables that we use comprise the fundamental variables examined in Campbell, Hilscher, and Szilagyi (2008) and are described in Online Appendix B.

The results of these regressions are presented in Table 4. We split the sample into financial and non-financial firms, defined by the firm's GICS sector from Compustat. Campbell, Hilscher, and Szilagyi (2008) consider only non-financial firms, and ratios such as leverage and book-to-market ratio are likely to be very different for financial firms than non-financial firms.

The results for non-financial firms indicate that the relations between CDS-implied and option-implied default probabilities and firm characteristics are similar. The implied default probabilities for non-financial firms are statistically significantly decreasing in profitability, book-to-market, and

[9] Results are qualitatively unchanged if we sample the last weekly observation of the quarter for the market variables and default probabilities, rather than using all weekly observations.

relative size when default probabilities are measured using either options or CDS. Both the default probabilities are statistically significantly increasing in leverage and return volatility. All of these relations are consistent with the results in Campbell, Hilscher, and Szilagyi (2008). For example, the negative association between book-to-market ratios and default probabilities is consistent with the positive relation between market-to-book equity ratio and bankruptcy in Campbell, Hilscher, and Szilagyi (2008). This result is also consistent with the evidence in Hovakimian, Kayhan, and Titman (2012), who find that firms with higher proportions of tangible assets (or low book-to-market ratios) have lower default probabilities. Additionally, both default probabilities are positively and statistically significantly associated with cash holdings, suggesting a greater precautionary saving motive for holding cash as in Acharya, Davydenko, and Strebulaev (2012). For only one variable, excess returns, does the relation to the two measures of default probability change signs in this sub-sample.

Results for financial firms are broadly similar to those for non-financial firms. When default probabilities are measured by options, the regression coefficient associated with firm characteristics are similar for financial and non-financial firms in terms of sign and statistical significance. We obtain similar findings for CDS implied default probabilities, with the exception of its relation to book-to-market. Overall, the results of this analysis suggest that the relation of firm characteristics to both CDS and option-implied default probabilities are consistent with the underlying fundamentals of the firm.

4 Sources of Difference Between Option- and CDS-Implied Default Probabilities

The evidence presented so far suggests that default probabilities inferred from CDS and option prices contain broadly similar information regarding the time-series properties of probabilities of default. Additionally, the evidence suggests that credit ratings and firm characteristics that are hypothesized to be related to the probability of default are related to both option- and CDS-implied default probabilities; that is, both estimates of default probabilities are capturing cross-sectional information.

However, as noted above, default probabilities estimated from the CDS and equity options market are not perfectly correlated. It is possible that these differences may simply arise from estimation error; in particular, the option-implied default probabilities are based on estimation of risk neutral moments and the imposition of the NIG distributional assumption. In this section, we consider various reasons that the probabilities from the two markets might differ.

The first possibility we consider is that differences in option and CDS prices are due to variation in rates of LGD. A second possibility is that aggregate and security-level liquidity may impact option and CDS prices, and therefore the imputed default probabilities derived from these markets. For example, at the security level, our option-based probability estimates begin by calculating implied volatility surfaces using out-of-the-money puts and calls. These contracts, especially deep out-of-the-money contracts, are likely to have less liquidity than near-the-money puts and calls. Further, we are interpolating the volatility surface at a maturity of 365 days, where there are likely to be fewer available contracts and lower liquidity. In addition, CDS are also likely to suffer from liquidity issues. We are using one-year CDS contracts, which have lower liquidity than five-year contracts. Finally, CDS are relatively sparsely traded at the beginning of our sample period and some contracts also suffer from liquidity issues during the financial crisis.

A third possibility involves the default threshold, α, that is used when estimating default probabilities from the options market. That is, higher α's may result in option-implied default probabilities that better match CDS-implied default probabilities in times when CDS-implied probabilities are high and for firms with poorer credit ratings. Thus, as in Chen, Collins-Dufresne and Goldstein (2008), it may be that the market perceives the default threshold for equity as being different during times of financial market stress or when a firm is closer to its default boundary; although our estimation method lets α vary as the credit rating varies, the market‘s perception of the default threshold may change at a higher frequency than a firm‘s credit rating. In addition, there may be other economic rationales for a varying default threshold, such as strategic bankruptcies; that is, in some circumstances, firms may find it beneficial to default even if they are solvent and able to make debt payments, as in Davydenko and Strebulaev (2007).

4.1 A Simple Measure of LGD

To analyze differences in the default probability estimates in these two markets, and in particular to explore whether these differences may be related to losses given default, we consider a simplified version of the relation between the CDS spread, the default probability, and the loss given default. If we consider a simple one-period CDS contract:

$$S_t = \left(Q_t^C\right) \cdot \left(LGD_t\right),$$

the CDS spread is the product of the default probability and the loss given default (LGD). While the relationship between spreads, default probabilities, and LGD is more precisely given in equation (3), this approximation provides useful intuition for understanding the relation between spreads, default probabilities and rates of loss given default.

Under the standard practice of assuming a constant rate of LGD of 60%, the spread and the CDS implied default probability in this simple model are perfectly correlated by construction, both cross-sectionally and in the time series. However, if the option-implied default probability is a valid estimate of the true default probability, then the ratio of CDS spread and option-implied default probability should provide an (approximate) estimate of loss given default that is allowed to vary across firms and through time. Of course, this ratio will also capture effects related to liquidity, mis-specification of the default threshold, and other estimation errors. We begin by calculating this ratio, denoted as $L\hat{G}D$, and considering its properties below.

4.1.1 Cross-market inferences

In Figure 2, we present the time series of average estimates of $L\hat{G}D$ across all firms in our sample, where option-implied default probabilities are calculated using both a threshold based on credit ratings (labeled Rating-derived), and a constant default threshold of 0.15. To limit the losses given default to 100%, we set values of $L\hat{G}D$ that are above 100% to missing.

The figure shows that the average $L\hat{G}D$ measures obtained using the two default thresholds are highly correlated. In addition, the behavior of both measures of $L\hat{G}D$ through time is consistent with the interpretation that the measure is related to loss given default; note that the average $L\hat{G}D$ varies strongly with business conditions, consistent with the relation between recovery rates and market fundamentals documented in Jankowitsch, Nagler, and Subrahmanyam (2014). In particular, the variation in $L\hat{G}D$ implies that average losses given default are high in the early 2000s (at approximately 45%), followed by declines in LGD to values of approximately 10% to 15% during the economic recovery of the mid-2000s. Losses given default again rise sharply to 55% during the financial crisis of 2007-2009, and then gradually decline. The secondary increase in LGD in 2011 is contemporaneous with the downgrade of U.S. debt in 2011. Over the recent period of 2013-2017, the losses given default hover around 15%. Note that the average LGD in economic expansions is substantially lower than the typical assumed (constant) rate of 60%, although it is consistent with the average historical LGD figures reported for large corporate borrowers (see, e.g., Global Credit Data (2018)).

Regardless of default thresholds, the evidence in Figure 2 that $L\hat{G}D$ varies with economic conditions is consistent with an inference that combining option-implied default probability measures with CDS data provides information about recovery rates, and indicates that the assumption that recovery rates are constant through time is a poor fit to the data. In addition, if both default probabilities and losses given default are countercyclical, note that estimates of default probability taken from CDS under the assumption that LGD are *constant* would result in default probability estimates that are more variable, and more right-skewed, than estimates taken from option prices, consistent with the evidence in Table 2; intuitively, if one assumes that LGD is constant instead of allowing LGD and default probabilities to covary positively, then estimated default probabilities must vary more in order to match the volatility in CDS rates. In the next sections, we examine the variation in $L\hat{G}D$ across sectors, while controlling for other factors such as liquidity.

4.1.2 Loss Given Default and Industry Sectors

We analyze variation in $L\hat{G}D$ across sectors. The results so far indicate that the cross-sectional variation in option-implied default probabilities is sensitive to the default threshold chosen, although the time-series information in $L\hat{G}D$ across different thresholds is similar. As a consequence, we examine LGD inferences for the two default thresholds presented in Figure 2: a constant $\alpha = 0.15$ and the firm-specific default thresholds based on credit rating (i.e., the Rating-derived α). We utilize the GICS sector definitions, which separate firms into 11 sectors; Energy, (EN) Materials (MA), Industrials (IN), Consumer Discretionary (CD), Consumer Staples (CS), Healthcare (HC), Financial (FI), Information Technology (IT), Telecommunications (TC), Utilities (UT) and Real Estate (RE). Sector classifications are obtained from Compustat.

In Table 5, we report the estimates obtained using both default thresholds. For each default threshold measure, we report the number of firms in the sample or sector in the first column, followed by the average $L\hat{G}D$ in the sample or sector; in the remaining columns, we present the fifth percentile and fiftieth percentile of the cross-sectional distribution of $L\hat{G}D$. This is followed by a column that reflects the percentile in the sample or sector where $L\hat{G}D$ exceeds 100%.

When default thresholds vary by credit rating, the median implied loss given default across the entire sample is 21%, or a recovery rate of 79%; note again that the median $L\hat{G}D$ using this approximation is lower than the typical Markit estimate of 60%. For nine of the eleven sectors, the median $L\hat{G}D$ ranges between 16% (Industrials) to 30% (Financials), with substantially higher $L\hat{G}D$ for Telecommunications and Real Estate. The number of firms with $L\hat{G}D$ greater than 1 is relatively small; there are six firms in total with these extreme $L\hat{G}D$ in four sectors. Overall, 99% of the firms have $L\hat{G}D$ lower than 1.0.

We obtain qualitatively similar results when the default threshold is held at a constant 0.15, with the range of $L\hat{G}D$ increasing in most cases. In the full sample, the median $L\hat{G}D$ is 0.32, and the range of $L\hat{G}D$ in nine out of eleven sectors is within the range of 16% to 44%.

For both default thresholds, we observe considerable variation in $L\hat{G}D$ across sectors, indicating that recovery rates vary cross-sectionally. Using both measures of default threshold, there are two sectors with higher $\hat{LGD}$: telecommunications and real estate, where the estimated median LGD increases to approximately 0.5 in the ratings-derived default threshold, and to values of median $\hat{LGD}$ at 98% and 68%, respectively, using the constant default threshold. Note that these two sectors also have a relatively small number of firm observations (10 or fewer); as a consequence, these results may be subject to greater estimation error. Overall, however, the correlation across sectors in the median estimate of $\hat{LGD}$ across the two default thresholds is quite high, at 0.97.

4.2 Liquidity, Default Probabilities, and Loss Given Default

While variation in LGD is one possible explanation for differences in risk neutral default probabilities across options and CDS, another possibility is that estimated probabilities in these two markets differ as a result of market frictions. As mentioned above, out-of-the-money options used to estimate risk-neutral moments and then option-implied default probabilities may be thinly traded; in addition, the liquidity of some CDS contracts is low. As a consequence, the marked increase in LGD around the crisis may reflect instead changes in market liquidity.[10] Note that the approximate relation between CDS rates, option-implied default probabilities and loss given default discussed in Section 4.1 implies that, in the absence of market frictions, the difference between the log CDS spread and the log default probability should be approximately equal to the log LGD. We use this approximation and estimate the extent to which variation in this difference is related to changes in various liquidity measures.

Illiquidity in the CDS and options markets may reflect both security-specific and market-wide variation in liquidity.[11] We are limited in measuring security-specific liquidity by the data available for options and CDS. In the case of options, we have information on bid-ask spreads, open interest, and volume. Since the default probabilities recovered from options are likely to depend most on

[10]It is possible that precipitous declines in market liquidity are associated with declines in asset values and thus increases in LGD (see, e.g., Brunnermeier and Pedersen (2009). If that is the case, our controls for market liquidity will cause the increases in loss given default during periods of market illiquidity to be estimated conservatively.

[11]Note that by using Fama and MacBeth (1973) regressions we are in effect including a time fixed effect in the regression. Thus, the results reported earlier supporting the interpretation of $\hat{LGD}$ as a proxy for loss given default are unlikely to be due to aggregate liquidity effects.

the prices of out-of-the-money options close to 365 days to maturity, we construct $SPREAD_t^O$, the average percentage bid-ask spread for the out-of-the-money options used in constructing our volatility surface. We also compute VOL_t^O and $OPEN_t^O$, the sum of volume and open interest for these contracts. In the case of CDS, we have a measure of the firm-specific depth for five-year CDS contracts, $DEPTH_t^C$. We assume that depth for the one-year contracts is correlated with the depth of the five-year contracts for each firm and use that as another measure of liquidity.

To capture aggregate liquidity, we use two measures from the fixed income security markets. First, we use the Treasury-Eurodollar spread, TED_t, measured as the difference in 90-day LIBOR and 90-day Treasury Bill yields. An increase in the TED spread can indicate an increase in interbank counterparty credit risk, and a consequent drop in funding liquidity. The second measure is the root mean squared error of the difference in market Treasury security yields from those implied by a Nelson-Siegel-Svensson model. This measure, $NOISE_t$, is investigated in Hu, Pan, and Wang (2013). The authors suggest that $NOISE_t$ is high when there is less arbitrage capital available in the Treasury market, a condition associated with lower liquidity. The TED spread is constructed using data from the Federal Reserve and $NOISE_t$ is obtained from Jun Pan's webpage.[12] Finally, we include a proxy for liquidity in the equity markets; Nagel (2012) suggests that a high level of the VIX index, VIX_t is associated with a high risk premium, and a consequent large reduction in liquidity provision, in equity markets. Data on the VIX are also obtained from the Federal Reserve.

We examine the extent to which liquidity effects influence the relation between CDS- and option-implied default probabilities by estimating the relation between changes in the (log) approximate loss given default and changes in the (log) liquidity variables, beginning at the aggregate level. That is, we estimate the parameters of a regression,

$$\begin{aligned}\Delta \hat{lgd}_{a,t} &= a_a + b_{a,1}\Delta ted_t + b_{a,2}\Delta noise_t + b_{a,3}\Delta vix_t + b_{a,4}\Delta spread_{a,t}^O \\ &\quad + b_{a,5}\Delta vol_{a,t}^O + b_{a,6}\Delta open_{a,t}^O + b_{a,7}\Delta depth_{a,t}^C + e_{a,t}, \qquad (6)\end{aligned}$$

where a indicates that we are measuring the quantity at the aggregate level, and aggregate variables

[12] We thank Jun Pan for making these data available at `http://www.mit.edu/~junpan/`.

are calculated as the cross-sectional average of individual time series observations. Lowercase variables are natural logs of their uppercase counterparts. The option-implied probability of default that is used to construct the estimate of loss given default uses a rating-derived default threshold. All variables are measured at the weekly horizon.

Results of this regression are reported in Table 6. When all liquidity variables are included, there is some evidence that changes in option market liquidity variables are associated with changes in CDS spreads. Specifically, the coefficient on $\Delta spread^O$ is significant at the 1% level, with an increase in spreads associated with a downward revision in the approximate loss given default. In contrast, changes in the VIX are associated with a significant increase in LGD, suggesting that some of the marked increase in LGD observed in the financial crisis may be associated with an increase in market-wide volatility. This result is consistent with the evidence in Nagel (2012), who shows that an increase in the VIX is associated with a reduction of equity arbitrage capital; alternatively, or in addition to a liquidity effect, the increase in market-wide volatility may be associated with a decline in asset values. Together, changes in these liquidity variables explain approximately 17% of the variation in changes in the log loss given default measure.

We also report the results of this regression across sectors in Table 7. The results are generally consistent with the results observed in the aggregate. Across all sectors with the exception of real estate, we continue to find evidence that changes in the VIX are positively associated with changes in the average $L\hat{G}D$ in the sector. Changes in option spreads are negatively and statistically significantly associated with changes in $L\hat{G}D$ in 6 out of the eleven sectors at the 5% or 10% critical level (specifically, consumer discretionary, consumer staples, healthcare, financial, information technology and utilities). Changes in option open interest are significantly negatively related to changes in $l\hat{g}d$ in the real estate sector. Finally, changes in the $NOISE$ measure are statistically significantly and positively related to changes in $l\hat{g}d$ in the energy and consumer staples sectors. In all of these cases, the R^2 measures indicate that liquidity measures account for less than 10% of the time series variation in changes in the approximate loss given default measure.

Using these regression results, we construct an alternative measure of approximate loss given default that controls for the effect of these liquidity measures. Specifically, we construct an alternative $L\hat{G}D$ by cumulating the residuals from equation (6). At time t, this liquidity-adjusted loss given default measure is calculated as

$$L\hat{G}D^{*}_{a,t} = \exp\left(\hat{a}_a + \sum_{j=0}^{t} \hat{e}_{a,t-j}\right),$$

where we initialize $L\hat{G}D^{*} = L\hat{G}D$ at time $t = 1$ (in January 2002) and at each period add the residual at time t given by the regression above. We exponentiate this series so that it can be compared to the $L\hat{G}D$ measure calculated previously.

Figure 3 presents the time series of this variable. The behavior of this liquidity-adjusted loss given default measure is broadly similar to the series plotted in Figure 2, with a relatively high implied loss given default during the early 2000s recession that declines in the mid 2000s, following by sharp increases in the financial crisis of 2008-2009. There are two noteworthy differences in the behavior of these loss given default measures. First, $L\hat{G}D^{*}$ is substantially higher than $L\hat{G}D$ beginning approximately in 2009. That is, adjusting for liquidity effects results in an estimate of loss given default $L\hat{G}D^{*}$ that is meaningfully higher during the financial crisis. Second, the liquidity-adjusted estimate of loss given default remains relatively elevated after the financial crisis, although it does decline, to approximately 0.35 in 2013, and then increases somewhat in the 2014-2017 period. Overall, estimates of loss given default that control for liquidity effects increase in economic downturns; indeed, our evidence suggests that liquidity-adjusted loss given default measures are more sensitive to economic states than LGD measures that do not control for liquidity.

5 Conclusion

In this paper, we propose a new method of estimating default probabilities for firms. Using option prices, we construct an estimate of the risk-neutral density; the default probability for the firm is the mass under the density up to the return that corresponds to a default event. We estimate default probabilities for different levels of default thresholds; we find that, although the level of estimated default probabilities is sensitive to the choice of default threshold, they are very highly correlated with one another and behave very similarly over time.

We examine the relationship of the option-implied default probabilities to default probabilities estimated from CDS prices, as well as their relation to firm characteristics and ratings categories. We find that option-implied default probabilities are strongly, but not perfectly, related to CDS default probabilities that assume constant losses given default, with the latter exhibiting higher variation and higher skewness. With respect to firm characteristics and ratings categories, the option-implied default probabilities behave as one would expect. Specifically, the default probabilities increase monotonically as ratings decline; in addition, default probabilities of non-financial firms are significantly and positively related to leverage and volatility; they are negatively and significantly related to profitability, book-to-market equity, and size.

If option-implied default probabilities are valid estimates of the firm's propensity of failure, then an examination of the relation between these probabilities and CDS prices should provide information about losses given default. We construct a simple measure of loss given default by calculating the ratio of CDS spreads and option-implied default probabilities. We find significant time-series variation in this ratio, which is related to economic conditions. While the ratio is highly correlated across default thresholds, we also find evidence of significant cross-sectional variation in this measure depending on sectors.

When we estimate the relation between CDS spreads and option-implied default probabilities, while controlling for liquidity effects, we find evidence that changes in the VIX is significantly and positively related to changes in log loss given default measures; we find weaker evidence that changes in the liquidity of the options and fixed income markets affect changes in LGD. Finally, after controlling for changes in liquidity effects, estimates of LGD inferred from CDS and option prices again show a strong relation to underlying business conditions. Overall, the equity option market may provide useful information with which to infer default probabilities, as well as the losses given default of underlying assets.

References

Acharya, Viral, Sergei A Davydenko, and Ilya A Strebulaev, 2012, Cash holdings and credit risk, *Review of Financial Studies* 25, 3572–3609.

Altman, Edward, B Bradi, A Resti, and A Sironi, 2005, The link between default and recovery rates: Theory, empirical evidence and implications, *Journal of Business* 78, 2203–2228.

Augustin, Patrick, and Fahad Saleh, 2017, A note on cds returns, unpublished manuscript, McGill University and New York University.

Augustin, Patrick, Marti G. Subrahmanyam, Dragon Yongjun Tang, and Sarah Qian Wang, 2016, Credit default swaps: Past, present, and future, forthcoming, *Annual Review of Financial Economics.*

Bakshi, Gurdip, N Kapadia, and Dilip Madan, 2003, Stock return characteristics, skew laws and the differential pricing of individual equity options, *Review of Financial Studies* 16, 101–143.

Bakshi, Gurdip, Dilip Madan, and Frank Zhang, 2006, Understanding the role of recovery in default risk models: Empirical comparisons and implied recovery rates, unpublished manuscript, University of Maryland.

Berndt, Antje, Rohan Douglas, Darrell Duffie, and Mark Ferguson, 2018, Corporate credit risk premia, unpublished manuscript,National Bureau of Economic Research.

Breeden, Douglas, and Robert Litzenberger, 1978, Prices of state contingent claims implicit in options prices, *Journal of Business* 51, 621–651.

Brunnermeier, Markus K, and Lasse Heje Pedersen, 2009, Market liquidity and funding liquidity, *Review of Financial Studies* 22, 2201–2238.

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

Capuano, Christian, 2008, The option-ipod: the probability of default implied by option prices based on entropy, IMF Working Paper 08-194.

Carr, Peter, and Liuren Wu, 2011, Simple robust linkages between american puts and credit protection, *Review of Financial Studies* 24, 473–505.

Chang, Bo Young, Peter Christoffersen, and Kris Jacobs, 2013, Market skewness risk and the cross-section of stock returns, *Journal of Financial Economics* 107, 46–68.

Chava, Sudheer, and Robert Jarrow, 2004, Bankruptcy prediction with industry effects, *Review of Finance* 8, 537–569.

Das, Sanjiv, and Raj Sundaram, 2007, An integrated model for hybrid securities, *Management Science* 53, 1439–1451.

Davydenko, Sergei A., and Ilya A. Strebulaev, 2007, Strategic actions and credit spreads: An empirical investigation, *Journal of Finance* 62, 2633–2671.

Dennis, Patrick, and Stuart Mayhew, 2002, Risk-neutral skewness: Evidence from stock options, *Journal of Financial and Quantitative Analysis* 37, 471–493.

Doshi, Hitesh, Redouane Elkamhi, and Chayawat Ornthanalai, 2018, The term structure of expected recovery rates, *Journal of Financial and Quantitative Analysis* 53, 2619–2661.

Duffie, Darrell, and Kenneth J Singleton, 1999, Modeling term structures of defaultable bonds, *Review of Financial Studies* 12, 687–720.

Eriksson, A, Eric Ghysels, and F Wang, 2009, The normal inverse gaussian distribution and the pricing of derivatives, *Journal of Derivatives* 16, 23–37.

Fama, Eugene F., and James D. MacBeth, 1973, Risk, return, and equilibrium: Empirical tests, *Journal of Political Economy* 81, 607–636.

Global Credit Data, 2018, LGD Report 2018 - Large Corporate Borrowers, Discussion paper, Global Credit Data Consortium.

Hansis, Alexandra, Christian Schlag, and Grigory Vilkov, 2010, The dynamics of risk-neutral implied moments: Evidence from individual options, unpublished manuscript, Goethe University Frankfurt.

Hovakimian, A., A. Kayhan, and S. Titman, 2012, Are corporate default probabilities consistent with the static trade-off theory?, *Review of Financial Studies* pp. 315–340.

Hu, Xing, Jun Pan, and Jiang Wang, 2013, Noise as information for illiquidity, *Journal of Finance* 68, 2341–2382.

Jankowitsch, Rainer, Florian Nagler, and Marti G Subrahmanyam, 2014, The determinants of recovery rates in the u.s. corporate bond market, *Journal of Financial Economics* 114, 155–177.

Le, Anh, 2015, Separating the components of default risk: A derivate-based approach, *The Quarterly Journal of Finance* 5, 1550005.

Madan, Dilip, and Haluk Unal, 1998, Pricing the risks of default, *Review of Derivatives Research* 2, 121–160.

——— , and Levent Güntay, 2003, Pricing the risk of recovery in default with absolute priority rule violation, *Journal of Banking and Finance* 27, 1001–1025.

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

Nagel, Stefan, 2012, Evaporating liquidity, *Review of Financial Studies* 25, 2005–2039.

Nelson, Charles R., and Andrew F. Siegel, 1987, Parsimonious modeling of yield curves, *Journal of Business* 60, 473–489.

O'Kane, Dominic, and Stuart Turnbull, 2003, Valuation of credit default swaps, *Lehman Brothers Fixed Income Quantitative Research Quarterly* pp. 1–18.

Schuermann, Til, 2004, What do we know about loss given default?, Wharton Financial Institutions Center Working Paper No. 04-01.

Svensson, Lars E. O., 1994, Estimating and interpreting forward interest rates: Sweden 1992-1994, NBER Working Paper 4871.

Vilsmeier, Johannes, 2014, Updating the option implied probability of default methodology, unpublished manuscript, Deutsche Bundesbank Discussion Paper No. 43/2014.

Table 1: Drop in Equity Prices Over 12 Months Prior to Bankruptcy

This table presents the magnitude of the drop in equity prices for firms that file bankruptcy over the previous 12 months. Bankruptcy filing dates are from the updated version of the data in Chava and Jarrow (2004) and are merged with CRSP data on equity prices and returns. We include only those firms that have CRSP data 12 months prior to the bankruptcy and in the month of bankruptcy filing and whose price declined over the 12 months prior to filing. The table presents summary statistics by credit rating for the decline in price by Standard and Poors credit rating for those firms for which ratings data are available. We present means, standard deviations, minima, and maxima of the ratio of the value of the stock price in the month of bankruptcy filing or delisting price to the price 12 months prior.

Rating	N	Mean	Std.	Min.	Max.
A	3	0.14	0.18	0.00	0.35
BBB	26	0.08	0.11	0.00	0.47
BB	67	0.13	0.20	0.00	0.98
B	197	0.12	0.16	0.00	0.95
CCC	59	0.20	0.22	0.00	0.84
CC	9	0.45	0.23	0.18	0.86
D	22	0.38	0.30	0.02	1.00

Table 2: Summary Statistics for Probabilities of Default

Table 2 presents summary statistics for risk neutral probabilities of default implied by credit default swap (CDS) spreads and prices of options on the equity of the same firm. CDS-implied default probabilities are calculated using a Nelson-Siegel-Svensson zero coupon term structure and constant one-year maturity CDS, assuming a loss given default of 60% on the underlying bond. Option-implied default probabilities are measured using the Bakshi, Kapadia, and Madan (2003) (BKM) procedure for computing risk neutral moments, then computing risk neutral probabilities of the Normal Inverse Gaussian distribution (NIG) on the basis of these moments. Risk neutral moments are computed using the implied volatility surface of options with 365 days to maturity. In Panel A, we present means, standard deviations, fifth, fiftieth, and ninety-fifth percentiles of the distribution of averages of probabilities implied by each security. In Panel B we present the fifth, fiftieth, and ninety-fifth percentiles of the distribution of correlations of a firm's option-implied and CDS-implied default probability. CDS data are from Markit and options data are from OptionMetrics. Data are sampled at the weekly frequency for 276 firms over the period January, 2002 through February, 2017.

Panel A: Distribution

	Mean	Std	p5	p50	p95
Option	2.87	2.61	0.71	2.03	7.77
CDS	1.98	3.75	0.15	0.77	6.56

Panel B: Correlations

p5	p50	p95
-0.10	0.52	0.88

Table 3: Summary Statistics for Probabilities of Default by Credit Rating

Table 3 presents summary statistics for risk neutral probabilities of default implied by credit default swap (CDS) spreads and prices of options on the equity of the same firm, grouped by credit rating. Firms with credit ratings augmented by '+' or '-' are grouped together; for example, rating 'BBB' refers to firms with a credit rating of 'BBB+', 'BBB', or 'BBB-'. CDS-implied default probabilities are calculated using a Nelson-Siegel-Svensson zero coupon term structure and constant one-year maturity CDS, assuming a loss given default of 60% on the underlying bond. Option-implied default probabilities are measured using the Bakshi, Kapadia, and Madan (2003) (BKM) procedure for computing risk neutral moments, then computing risk neutral probabilities of the Normal Inverse Gaussian distribution (NIG) on the basis of these moments. Risk neutral moments are computed using the implied volatility surface of options with 365 days to maturity. We calculate the average default probability for each firm, conditional on its ratings group, and report the mean and median of these averages by ratings group. CDS data are from Markit, options data are from OptionMetrics, and ratings data are from Compustat. Options and CDS data are sampled at the weekly frequency, and ratings at the monthly frequency for 276 firms over the period January, 2002 through February, 2017.

		Option		CDS	
		Mean	Median	Mean	Median
AAA	9	0.45	0.35	0.23	0.19
AA	32	0.90	0.87	0.44	0.24
A	116	1.35	1.37	0.61	0.32
BBB	158	2.16	2.08	1.27	0.72
BB	88	3.71	3.64	2.87	2.34
B	49	6.73	6.09	4.72	3.59
CCC+ and below	9	14.39	14.51	15.38	11.48

Table 4: Relation between Firm Characteristics and Default Probabilities

Table 4 examines the relationship between default probabilities implied by either CDS spreads using a 60% loss given default assumption or options with a critical threshold that varies with credit rating, and firm-specific characteristics. Default probabilities at the end of each month are regressed on a set of nine firm-specific variables: $NIMTA$, the ratio of net income to market value of total assets, $TLMTA$, the ratio of total liabilities to market value of assets, $EXRET$, the monthly log return on the firm's equity in excess of that of the S&P 500, $SIGMA$, the volatility of the firm's equity return over the past three months, $RSIZE$, the log ratio of the market capitalization of the firm's equity to that of the S&P 500, $CASHMTA$, and BM, the firm's ratio of book value of equity to market value of equity. Point estimates are the average of monthly regression coefficients, and standard errors in parentheses are corrected using the Newey-West procedure. We present results for financial firms and for firms excluding financial firms, defined as those in GICS sector 40. Data for CDS are obtained from Markit, data for options is obtained from Option Metrics, return information is obtained from CRSP, financial statement and ownership information is obtained from Compustat.

	CDS		Option	
	Financial	Non-Financial	Financial	Non-Financial
$NIMTA$	-17.32^{***}	-11.20^{***}	-24.37^{***}	-5.99^{***}
	(2.18)	(0.34)	(1.30)	(0.22)
$TLMTA$	2.90^{***}	1.54^{***}	1.89^{***}	2.33^{***}
	(0.13)	(0.04)	(0.08)	(0.02)
$EXRET$	-0.46^{***}	-0.35^{***}	1.56^{***}	0.38^{***}
	(0.15)	(0.04)	(0.09)	(0.02)
$SIGMA$	7.45^{***}	7.33^{***}	1.99^{***}	4.84^{***}
	(0.09)	(0.04)	(0.05)	(0.02)
$RSIZE$	-1.00^{***}	-0.80^{***}	-0.86^{***}	-0.36^{***}
	(0.02)	(0.01)	(0.01)	(0.00)
$CASHMTA$	4.51^{***}	1.62^{***}	3.91^{***}	2.21^{***}
	(0.29)	(0.09)	(0.18)	(0.06)
BM	0.04^{***}	-0.47^{***}	-0.06^{***}	-0.32^{***}
	(0.01)	(0.01)	(0.01)	(0.01)
R^2	0.54	0.46	0.40	0.59

*,**,*** denotes significance at the 10%, 5%, and 1% critical level, respectively.

Table 5: Summary Statistics for Ratio of CDS Spreads to Option Probabilities

Table 5 presents summary statistics for the ratio of one year CDS spreads and default probabilities implied by risk neutral probabilities of default as measured by options. Option-implied default probabilities are measured using the Bakshi, Kapadia, and Madan (2003) (BKM) procedure for computing risk neutral moments, then computing risk neutral probabilities of the Normal Inverse Gaussian distribution (NIG) on the basis of these moments. Risk neutral moments are computed using the implied volatility surface of options with 365 days to maturity. CDS data are from Markit and options data are from OptionMetrics. The table presents the number of firms, fifth and fiftieth percentiles of the cross-sectional distribution of average ratio of one year CDS spreads to option-implied default probabilities, and the percentile at which the ratio exceeds 1.0. Results are presented for all firms in the sample and with GICS sectors and two default thresholds; one based on Rating-derived thresholds and the other a constant threshold of 0.15. Data are sampled at the weekly frequency for 276 firms over the period January, 2002 through February, 2017.

		Rating-Derived α			$\alpha = 0.15$		
Sector	N	p5	p50	$L = 100$	p5	p50	$L = 100$
All	276	0.08	0.21	99	0.08	0.32	84
Energy	26	0.12	0.19	100	0.09	0.27	85
Materials	19	0.08	0.20	100	0.10	0.33	95
Industrials	37	0.08	0.16	95	0.07	0.16	84
Discretionary	50	0.11	0.24	100	0.12	0.44	78
Staples	26	0.10	0.20	100	0.09	0.18	81
Healthcare	30	0.10	0.19	100	0.07	0.27	97
Financials	32	0.14	0.30	94	0.10	0.34	91
Technology	24	0.08	0.27	100	0.07	0.41	83
Telecommunications	10	0.08	0.53	90	0.08	0.98	50
Utilities	16	0.11	0.20	100	0.13	0.32	75
Real Estate	6	0.17	0.47	83	0.27	0.68	83

Table 6: Liquidity, CDS Spreads and Option-Implied Default Probabilities

Table 6 presents the results of regressions of changes in the log approximate loss given default on changes in log liquidity variables. The regressions are specified as

$$\begin{aligned} \Delta \hat{lgd}_{a,t} &= a_a + b_{a,1}\Delta ted_t + b_{a,2}\Delta noise_t + b_{a,3}\Delta vix_t + b_{a,4}\Delta spread^O_{a,t} \\ &\quad + b_{a,5}\Delta vol^O_{a,t} + b_{a,6}\Delta open^O_{a,t} + b_{a,7}\Delta depth^C_{a,t} + e_{a,t}, \end{aligned}$$

where $\hat{lgd}_{a,t}$ is the log of an approximate loss given default measure and is constructed by subtracting the log of the option-implied default probability from the log of the one-year CDS spread. ted_{t+1} is the log TED spread, the difference between the yield on 90-day LIBOR and 90-day Treasury Bills, $noise_{t+1}$ is the log of the noise measure from Hu, Pan, and Wang (2013), vix_{t+1} is the log VIX index, $vol^O_{i,t+1}$ is the log of the sum of volume for out-of-the-money (OTM) options on firm i, $open^O_{i,t+1}$ is the log sum of open interest on firm i's OTM options, $spread^O_{i,t+1}$ is the log of the average percentage bid-ask spread for firm i's OTM options, and $depth^C_{i,t+1}$ is the depth of 5-year CDS contracts for firm i. Options data is from OptionMetrics, CDS data from Markit, financial market data from the Federal Reserve, and the noise measure from Jun Pan's website. The table first reports results of regressions of the change in log loss given default on the explanatory variables, where the aggregate is constructed as the cross-sectional average of each week's observations. The table also reports results for the fifth, twenty-fifth, median, seventy-fifth, and ninety-fifth percentile coefficient estimates and their associated standard errors for firm-specific regressions. The sample covers 276 firms over the period January, 2002 through February, 2017.

	Δted	$\Delta noise$	Δvix	Δvol^O	$\Delta open^O$	$\Delta spread^O$	$\Delta depth^C$	R^2
Aggregate	0.012	0.024	0.152	0.003	0.033	-0.113	-0.027	0.168
SE	(0.019)	(0.016)	(0.015)	(0.006)	(0.030)	(0.029)	(0.033)	
p5	-0.493	-0.215	-0.195	-0.108	-0.724	-0.525	-0.133	0.012
SE	(0.371)	(0.596)	(0.147)	(0.132)	(0.517)	(0.242)	(0.053)	
p25	-0.082	-0.022	0.029	-0.034	-0.133	-0.215	-0.023	0.029
SE	(0.174)	(0.055)	(0.065)	(0.031)	(0.406)	(0.122)	(0.023)	
p50	0.007	0.040	0.165	-0.008	-0.002	-0.079	0.008	0.050
SE	(0.080)	(0.045)	(0.079)	(0.026)	(0.164)	(0.093)	(0.017)	
p75	0.082	0.115	0.259	0.016	0.110	0.031	0.039	0.086
SE	(0.041)	(0.057)	(0.073)	(0.038)	(0.258)	(0.274)	(0.054)	
p95	0.285	0.330	0.457	0.082	0.447	0.446	0.122	0.256
SE	(0.098)	(0.170)	(0.064)	(0.059)	(0.183)	(0.393)	(0.090)	

Table 7: Liquidity, CDS Spreads and Option-Implied Default Probabilities by Sector

Table 7 presents the results of regressions of changes in the log approximate loss given default on changes in log liquidity variables. The regressions are specified as

$$\begin{aligned} \Delta \hat{lgd}_{a,t} &= a_a + b_{a,1}\Delta ted_t + b_{a,2}\Delta noise_t + b_{a,3}\Delta vix_t + b_{a,4}\Delta spread^O_{a,t} \\ &\quad + b_{a,5}\Delta vol^O_{a,t} + b_{a,6}\Delta open^O_{a,t} + b_{a,7}\Delta depth^C_{a,t} + e_{a,t}, \end{aligned}$$

where $\hat{lgd}_{a,t}$ is the log of an approximate loss given default measure and is constructed by subtracting the log of the option-implied default probability from the log of the one-year CDS spread. ted_{t+1} is the log TED spread, the difference between the yield on 90-day LIBOR and 90-day Treasury Bills, $noise_{t+1}$ is the log of the noise measure from Hu, Pan, and Wang (2013), vix_{t+1} is the log VIX index, $vol^O_{i,t+1}$ is the log of the sum of volume for out-of-the-money (OTM) options on firm i, $open^O_{i,t+1}$ is the log sum of open interest on firm i's OTM options, $spread^O_{i,t+1}$ is the log of the average percentage bid-ask spread for firm i's OTM options, and $depth^C_{i,t+1}$ is the depth of 5-year CDS contracts for firm i. Options data is from OptionMetrics, CDS data from Markit, financial market data from the Federal Reserve, and the noise measure from Jun Pan's website. The table reports results for series aggregated across two-digit GICS codes, Energy (EN), Materials (MA), Industrials (IN), Consumer Discretionary (CD), Consumer Staples (CS), Healthcare (HC), Financials (FI), Information Technology (IT), Telecommunications (TC), Utilities (UT) and Real Estate (RE). Data are sampled at the weekly frequency from January 2002 through February 2017.

	Δted	$\Delta noise$	Δvix	Δvol^O	$\Delta open^O$	$\Delta spread^O$	$\Delta depth^C$	$\bar{R}^2$
Energy	−0.005	0.065	0.218	−0.010	−0.070	0.001	0.009	0.075
	(0.038)	(0.032)	(0.031)	(0.013)	(0.062)	(0.059)	(0.037)	
Materials	0.011	0.011	0.192	0.006	−0.015	−0.014	−0.042	0.041
	(0.045)	(0.038)	(0.037)	(0.015)	(0.074)	(0.071)	(0.036)	
Industrials	0.041	−0.046	0.081	0.018	−0.056	−0.075	−0.017	0.022
	(0.041)	(0.035)	(0.034)	(0.014)	(0.068)	(0.065)	(0.046)	
Cons. Disc.	0.020	0.012	0.174	0.003	0.058	−0.129	−0.006	0.068
	(0.035)	(0.029)	(0.028)	(0.012)	(0.057)	(0.054)	(0.042)	
Cons. Staples	−0.006	0.083	0.096	0.027	0.042	−0.144	0.046	0.045
	(0.037)	(0.031)	(0.030)	(0.013)	(0.061)	(0.058)	(0.041)	
Healthcare	−0.025	0.012	0.132	−0.009	0.039	−0.126	−0.087	0.056
	(0.034)	(0.028)	(0.028)	(0.012)	(0.056)	(0.053)	(0.039)	
Financials	−0.002	0.053	0.155	−0.003	0.098	−0.158	−0.023	0.062
	(0.036)	(0.030)	(0.029)	(0.012)	(0.059)	(0.056)	(0.034)	
Info. Tech.	0.047	0.007	0.048	−0.019	0.158	−0.176	−0.077	0.014
	(0.058)	(0.051)	(0.047)	(0.021)	(0.096)	(0.092)	(0.051)	
Telecom.	0.023	0.115	0.166	0.029	0.111	−0.218	−0.141	0.024
	(0.101)	(0.088)	(0.083)	(0.036)	(0.168)	(0.165)	(0.060)	
Utilities	0.015	0.077	0.175	−0.020	0.176	−0.182	−0.074	0.033
	(0.065)	(0.055)	(0.054)	(0.022)	(0.110)	(0.109)	(0.057)	
Real Estate	−0.048	0.001	0.154	−0.002	−0.460	0.147	0.019	0.012
	(0.122)	(0.106)	(0.103)	(0.043)	(0.208)	(0.211)	(0.075)	

Figure 1: Implied Default Probabilities

Figure 1 plots the time series of aggregate default probabilities. Default probabilities are measured using one year CDS spreads with an assumed rate of loss given default of 60% and using the risk neutral distribution implied by option prices with a ratings-dependent default threshold and a constant threshold $\alpha = 0.15$. CDS data are obtained from Markit and options data from OptionMetrics. Data are sampled at the weekly frequency and aggregated by taking the average across the firms in the sample. The data cover the period January 2002 through February 2017, and cover 276 firms.

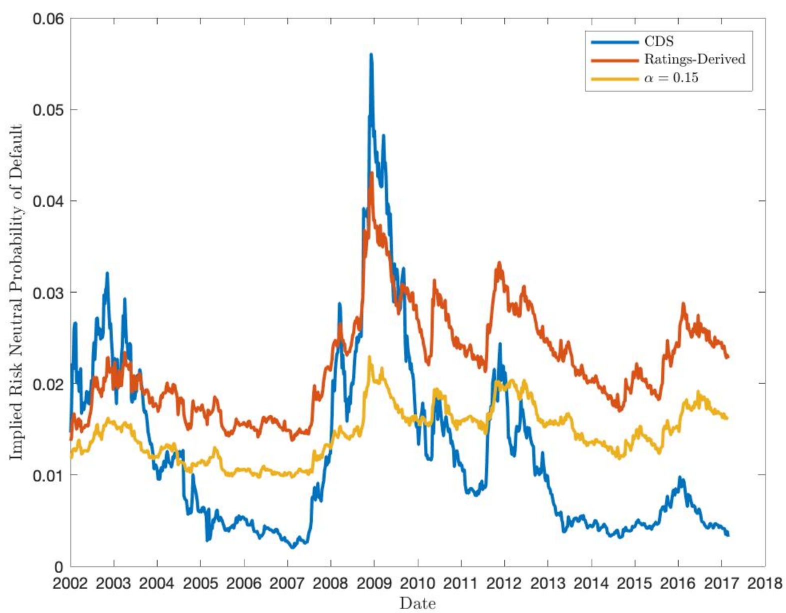

Figure 2: Differences in log CDS Spreads and Option-Implied Default Probabilities

Figure 2 plots the time series of the ratio of one-year CDS spreads and option-implied default probabilities aggregated over firms, denoted as $L\hat{G}D$. Option-implied default probabilities are calculated with a rating-derived default threshold and a constant threshold $\alpha = 0.15$. Options data are from OptionMetrics and CDS data are from Markit. Data cover 276 firms over the period January 2002 through February 2017.

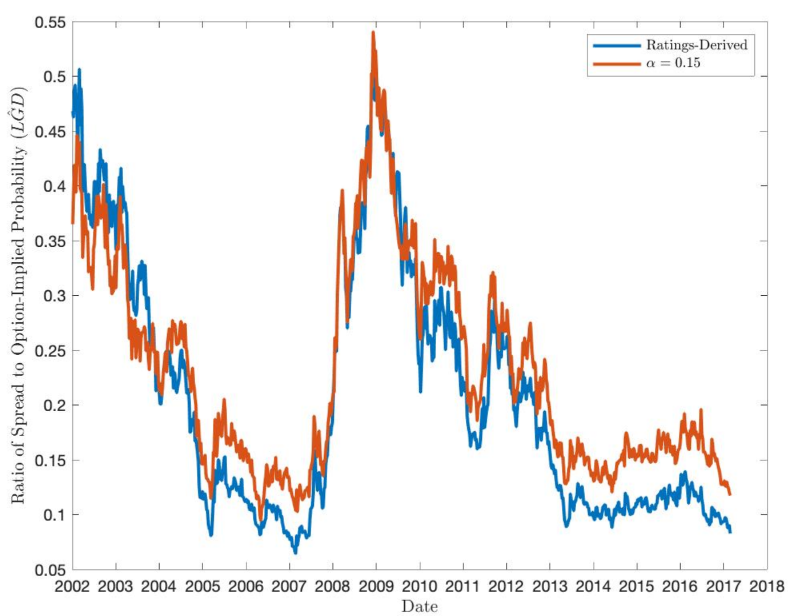

Figure 3: Liquidity-adjusted Loss Given Default

Figure 3 presents a liquidity-adjusted measure of the loss given default, represented by the cumulative residual from a regression of the change in approximate log loss given default on a set of variables measuring liquidity. Specifically, we regress changes in the log aggregate ratio of one year CDS spread to option-implied default probability on changes in the log of the aggregate of seven measures of liquidity; the TED spread, ted, the noise measure from Hu, Pan, and Wang (2013), $noise$, the VIX index, vix, the sum of OTM option volume, vol^O, the sum of OTM option open interest, $open^O$, the average OTM option bid-ask spread, $spread^O$, and the depth of 5-year CDS contracts, and $depth^C$:

$$\begin{aligned} \hat{lgd}_t &= a + b_1 \Delta ted_t + b_2 \Delta noise_t + b_3 \Delta vix_t + b_4 \Delta vol_t^O \\ &\quad + b_5 \Delta open_t^O + b_6 \Delta spread_t^O + b_7 \Delta depth_t^C + e_{t+1}, \end{aligned}$$

where aggregates of the variables vol^O, $open^O$, $spread^O$, and $depth^C$ are cross-sectional averages or sums of firm-level variables. The loss given default net of liquidity is

$$\hat{LGD}^* = \exp\left(\hat{lgd}_1 + \hat{a} + \sum_{j=0}^{t-1} \hat{e}_{it-j} \right).$$

Data for the construction of the TED spread and the VIX are obtained from the Federal Reserve. The noise measure is taken from Jun Pan's website. Data on option volume, open interest, and bid-ask spreads are from OptionMetrics. CDS contract depth is obtained from Markit. Data are sampled at the weekly frequency over the period January 2002 through February 2017.