P Feldhütter and S M Schaefer
The Myth of the Credit Spread Puzzle
Article

Feldhütter, P and Schaefer, S M

(2018)

The Myth of the Credit Spread Puzzle.

Review of Financial Studies, 31 (8). pp. 2897-2942. ISSN 0893-9454

DOI: https://doi.org/10.1093/rfs/hhy032

Oxford University Press (OUP)
https://academic.oup.com/rfs/article/31/8/2897/495...

The Myth of the Credit Spread Puzzle

Peter Feldhütter
Copenhagen Business School

Stephen M. Schaefer
London Business School

Are standard structural models able to explain credit spreads on corporate bonds? In contrast to much of the literature, we find that the Black-Cox model matches the level of investment-grade spreads well. Model spreads for speculative-grade debt are too low, and we find that bond illiquidity contributes to this underpricing. Our analysis makes use of a new approach for calibrating the model to historical default rates that leads to more precise estimates of investment-grade default probabilities. (*JEL* C23, G12, G13)

Received October 25, 2016; editorial decision January 12, 2018 by Editor Andrew Karolyi. Authors have furnished an Internet Appendix, which is available on the Oxford University Press Web site next to the link to the final published paper online.

The structural approach to credit risk, pioneered by Merton (1974) and others, represents the leading theoretical framework for studying corporate default risk and pricing corporate debt. While the models are intuitive and simple, many studies find that, once calibrated to match historical default and recovery rates and the equity premium, they fail to explain the level of actual investment-grade credit spreads, a result referred to as the "credit spread puzzle."

We are grateful to the Q-Group for the Jack Treynor Prize. We are also grateful for very valuable comments from the editor and an anonymous referee and suggestions received from Antje Berndt (discussant), Harjoat Bhamra (discussant), Hui Chen (discussant), Mike Chernov, Darrell Duffie, Stefano Giglio (discussant), Jean Helwege (discussant), Ralph Koijen, David Lando, Erik Loualiche (discussant), Gustavo Manso (discussant), Robert Merton, Jens-Dick Nielsen, Lasse Heje Pedersen, Scott Richardson, and Batchimeg Sambalaibat (discussant) and seminar participants at the Q-Group Spring 2016 meeting, NBER Asset Pricing 2015 conference, Utah Winter Finance Conference 2015, EFA 2015, HKUST Finance Symposium 2014, WFA 2014, SFS Cavalcade 2014, ESSFM 2014 in Gerzensee, Fifth Risk Management Conference 2014, UNC's Roundtable 2013, UCLA Anderson School of Management, UC San Diego, Financial Conduct Authority, AQR, London Business School, Cass Business School, Vienna Graduate School of Finance, Stockholm School of Economics, Tilburg University, Duisenburg School of Finance Amsterdam, NHH Bergen, Rotterdam School of Management, University of Southern Denmark, Copenhagen Business School, Brunel University, Surrey Business School, Henley Business School, Bank of England, and University of Cambridge. We are particularly grateful for the research assistance provided by Ishita Sen. Support from the Center for Financial Frictions (FRIC) [grant no. DNRF102] is gratefully acknowledged. Supplementary data can be found on *The Review of Financial Studies* Web site.
Send correspondence to Peter Feldhütter, Copenhagen Business School, Solbjerg Plads 3, 2000 Frederiksberg, Denmark; telephone: +45 38153753. E-mail: pf.fi@cbs.dk.

doi:10.1093/rfs/hhy032 Advance Access publication March 24, 2018

Papers that find a credit spread puzzle typically use Moody's historical default rates, measured over a period of around 30 years and starting from 1970, as an estimate of the expected default rate.[1] Our starting point is to show that the appearance of a credit spread puzzle strongly depends on the period over which historical default rates are measured. For example, Chen, Collin-Dufresne, and Goldstein (2009) use default rates from 1970 to 2001 and find BBB-AAA model spreads of 57–79 basis points (bps) (depending on maturity), values that are substantially lower than historical spreads of 94–102 bps. If, instead, we use Moody's default rates for 1920–2001, model spreads are 91–112 bps, a range that is in line with historical spreads.

Using simulations, we demonstrate two key points about historical default rates. The first is, over sample periods of around 30 years that are typically used in the literature, there is a large sampling error in the observed average rate. For example, if the true 10-year BBB cumulative default probability were 5.09%,[2] a 95% confidence band for the realized default rate measured over 31 years would be $[1.15\%, 12.78\%]$. Intuitively, the large sample error arises because defaults are correlated and 31 years of data only give rise to three nonoverlapping 10-year intervals. As a result of the large sampling error, when historical default rates are used as estimates of ex ante default probabilities, the difference between actual spreads and model spreads needs to be large—much larger, for example, than that found for the BBB-AAA spread mentioned above—to be interpreted as statistically significant evidence against the model.

Second, and equally crucial, distributions of average historical investment-grade default rates are highly positively skewed. Most of the time we see few defaults, but, occasionally, we see many defaults, meaning that there is a high probability of observing a rate that is below the actual mean. Positive skewness is likely to lead to the conclusion that a structural model underpredicts investment-grade spreads even if the model is correct. The reason for the presence of skewness is that defaults are correlated across firms as a result of the common dependence of individual firm values on systematic ("market") shocks. To see why correlation leads to skewness, we can think of a large number of firms with a default probability (over some period) of 5% and where their defaults are perfectly correlated. In this case we will observe a zero default rate 95% of the time (and a 100% default rate 5% of the time), and so the realized default rate will underestimate the default probability 95% of the time. If the average default rate is calculated over three independent periods, the realized default rate will still underestimate the default probability $0.95^3 = 85.74\%$ of the time.

We propose a new approach to estimate default probabilities. Instead of using the historical default rate at a single maturity and rating as an estimate of the default probability for this same maturity and rating, we use a wide

[1] See, for example, Leland (2006), Cremers, Driessen, and Maenhout (2008), Chen, Collin-Dufresne, and Goldstein (2009), and Huang and Huang (2012).

[2] This is the number reported by Moody's for 1970–2001.

cross-section of default rates at different maturities and ratings. We use the Black and Cox (1976) model and what ties default probabilities for firms with different ratings together in the model is that we assume that they will, nonetheless, have the same default boundary. (The default boundary is the value of the firm, measured as a fraction of the face value of debt, below which the firm defaults.) This is reasonable since, if the firm were to default, there is no obvious reason the default boundary would depend on the rating the firm had held previously.

We show in simulations that our approach results in much more precise and less skewed estimates of investment-grade default probabilities. For the estimated 10-year BBB default probability, for example, the standard deviation and skewness using the new approach are only 16% and 4%, respectively, of those using the existing approach. The improved precision is partly the result of the fact that we combine information across 20 maturities and 7 ratings and default probability estimates from different rating/maturity pairs are imperfectly correlated. But, to a significant extent, it is the result of combining default information on investment-grade and high-yield defaults. Because defaults occur much more frequently in high-yield debt, these firms provide more information on the location of the default boundary. Since the boundary is common to investment-grade and high-yield debt, when we combine investment-grade and high-yield default data, we "import" the information on the location of the default boundary from high-yield to investment-grade debt. The reduction in skewness is also the result of including default rates that are significantly higher than those for BBB debt. While a low default rate for investment-grade debt produces a positive skew in the distribution of defaults, a default rate of 50% produces a symmetric distribution and, for even higher default rates, the skew is actually negative.

We use our estimation approach and the Black-Cox model to investigate spreads over the period 1987–2012. Our data set consists of 256,698 corporate bond yield spreads to the swap rate of noncallable bonds issued by industrial firms and is more extensive than those previously used in the literature. Our implementation of the Black-Cox model is new to the literature in that it allows for cross-sectional and time-series variation in firm leverage and payout rate while matching historical default rates. Applying our proposed estimation approach, we estimate the default boundary such that average model-implied default probabilities match average historical default rates from 1920 to 2012. In calibrating the default boundary we use a constant Sharpe ratio and match the equity premium, but, once we have implied out the single firm-wide default boundary parameter, we compute firm- and time-specific spreads using standard "risk-neutral" pricing formulae.

We first explore the difference between average spreads in the Black-Cox model and actual spreads. The average model spread across all investment-grade bonds with a maturity between 3 and 20 years is 111 bps, whereas the average actual spread is 92 bps. A confidence band for the model spread that takes into account uncertainty in default probabilities is [88 bps; 128 bps]; thus there is no statistical difference between actual and model investment spreads. For speculative-grade bonds, the average model spread is 382 bps, whereas the actual spread is 544 bps, and here the difference is statistically highly significant. We also sort bonds by the actual spread and find that actual and model-implied spreads are similar, except for bonds with a spread of more than 1,000 bps. For example, for bonds with an actual spread between 100–150 bps the average actual spread is 136 bps, whereas the average model-implied spread is 121 bps. Importantly, the results are similar if we calibrate the model using default rates from 1970 to 2012 rather than from 1920 to 2012, thus resolving the problem described above that results in the earlier literature depend significantly on the historical period chosen to benchmark the model.

To study the time series, we calculate average spreads on a monthly basis and find that for investment-grade bonds there is a high correlation of 93% between average actual spreads and model spreads. Note that the model-implied spreads are "out-of-sample" predictions in the sense that actual spreads are not used in the calibration. Furthermore, for a given firm only changes in leverage and the payout rate—calculated using accounting data and equity values—lead to changes in the firm's credit spread. For speculative-grade bonds the correlation is only 40%, showing that the model has a much harder time matching spreads for low-quality firms.

Although *average* investment-grade spreads are captured well on a monthly basis, the model does less well at the individual bond level. Regressing individual investment-grade spreads on those implied by the Black-Cox model gives an R^2 of only 44%, so at the individual bond level less than half the variation in investment-grade spreads is explained by the model. For speculative-grade spreads the corresponding R^2 is only 13%.

We also investigate the potential contribution of bond illiquidity to credit spreads. We use bond age as the liquidity measure and double sort bonds on liquidity and credit quality. For investment-grade bonds we find no relation between bond liquidity and spreads, consistent with the ability of the model to match actual spreads and the finding in Dick-Nielsen, Feldhütter, and Lando (2012) that outside the 2007–2008 financial crisis illiquidity premiums in investment-grade bonds were negligible. For speculative-grade bonds we find a strong relation between bond liquidity and yield spreads, suggesting that bond liquidity may explain much of the underpricing of speculative-grade bonds.

In this paper we use the Black and Cox (1976) model as a lens through which to study the credit spread puzzle. The results in Huang and Huang (2012) show that many structural models which appear very different in fact generate similar spreads once the models are calibrated to the same default probabilities, recovery rates, and the equity premium. The models tested in Huang and Huang (2012) include features such as stochastic interest rates, endogenous default, stationary leverage ratios, strategic default, time-varying asset risk premiums, and jumps in the firm value process, yet all generate a similar level of credit spread. To the extent that different structural models produce similar investment-grade default probabilities under our estimation approach, our finding that the Black-Cox model matches average investment-grade spreads is likely to hold for a wide range of structural models.

An extensive literature tests structural models. Leland (2006), Cremers, Driessen, and Maenhout (2008), Chen, Collin-Dufresne, and Goldstein (2009), Chen (2010), Huang and Huang (2012), Chen, Cui, He, and Milbradt (2017), Bai (2016), Bhamra, Kuehn, and Strebulaev (2010), and Zhang, Zhou, and Zhu (2009) use the historical default rate at a given rating and maturity to estimate the default probability at that maturity and rating. We show that this test is statistically weak. Eom, Helwege, and Huang (2004), Ericsson, Reneby, and Wang (2015), and Bao (2009) allow for heterogeneity in firms and variation in leverage ratios, but do not calibrate to historical default rates. Bhamra, Kuehn, and Strebulaev (2010) observe that default rates are noisy estimators of default probabilities, but do not propose a solution to this problem as we do.

1. A Motivating Example

There is a tradition in the credit risk literature of using Moody's average realized default rate for a given rating and maturity as a proxy for the corresponding ex ante default probability. This section provides an example showing that the apparent existence or nonexistence of a credit spread puzzle depends on the particular period over which the historical default rate is measured. Later in the paper we describe an alternative approach for extracting default probability estimates from historical default rates that not only provides much greater precision but is also less sensitive to the sample period chosen.

To understand how Moody's calculates default frequencies, consider the 10-year BBB cumulative default frequency of 5.09% used in Chen, Collin-Dufresne, and Goldstein (2009).[3] This number is published in Moody's (2002) and is based on default data for the period 1970–2001. For the year 1970, Moody's identifies a cohort of BBB-rated firms and then records how many of these default over the next 10 years. The 10-year BBB default frequency for 1970 is the number of defaulted firms divided by the number in the 1970 cohort. The average default rate of 5.09% is calculated as the average of the twenty-two 10-year default rates for the cohorts formed at yearly intervals over the period 1970–1991.

A large part of the literature has focused on the BBB-AAA spread at 4- and 10-year maturities. In our main empirical analysis (Section 3), we study a much wider range of ratings and maturities but for now, to keep our example simple, we also focus on the BBB-AAA spread. For a given sample period we use the BBB and AAA average default rates for the 4- and 10-year

[3] Moody's reports a 10-year BBB default rate of 5.09% (Exhibit 32), and Chen, Collin-Dufresne, and Goldstein (2009) use 4.89%. We use Moody's reported number.

horizons reported by Moody's. Following the literature (e.g., Chen et al. 2009; Huang and Huang 2012; and others) we first benchmark a model to match these default rates, one at a time. Using the benchmarked parameters we then compute risk-neutral default probabilities and, from these, credit spreads. Following Eom, Helwege, and Huang (2004), Bao (2009), Huang and Huang (2012), and others, we assume that if default occurs, investors receive (at maturity) a fraction of the originally promised face value, but now with certainty. The credit spread, s, is then calculated as

$$s = y - r = -\frac{1}{T}\log[1-(1-R)\pi^{Q}(T)], \tag{1}$$

where R is the recovery rate, T is the bond maturity, and $\pi^{Q}(T)$ is the risk-neutral default probability. Throughout our analysis we employ the Black-Cox model (Black and Cox 1976). Appendix A provides the model details.

We use our average parameter values for the period 1987–2012 estimated in Section 3 and Chen, Collin-Dufresne, and Goldstein's (2009) estimates of the Sharpe ratio and recovery rate. We estimate the default boundary by matching an observed default frequency. The default boundary is the value of the firm, measured as a fraction of the face value of debt, below which the firm defaults. Following Chen, Collin-Dufresne, and Goldstein (2009) and others, we carry this out separately for each maturity and rating such that, conditional on the other parameters, the model default probability matches the reported Moody's default frequency. For each maturity and rating we then use the benchmarked default boundary and calculate the credit spread using Equation (1).

The solid bars in Figure 1 show estimates of the actual BBB-AAA corporate bond credit spread from a number of papers. For both the 4- and 10-year maturities, the estimated BBB-AAA spread is in the range of 98–109 bps with the notable exception of Huang and Huang's (2012) estimate of the 10-year BBB-AAA of 131 bps. (Huang and Huang use both callable and noncallable bonds in their estimate of the spread and this may explain why it is higher.)

Using Moody's average default rates from the period 1970–2001, the 4- and 10-year BBB-AAA spreads in the Black-Cox model are 52 and 72 bps, respectively. These model estimates are substantially below actual spreads, a finding that has been coined the "credit spread puzzle."

Figure 1 also shows the model-implied spreads using Moody's average historical default rates from 1920 to 2002 (default rates from 1920 to 2001 are not available). Using default rates from this longer period, the model-implied spreads are substantially higher: the 4- and 10-year BBB-AAA spreads are 87 and 104 bps, respectively. Thus, when we use default rates from a longer time period the puzzle largely disappears.

To emphasize that this conclusion is not specific to the Black-Cox model, Figure 1 also shows the four spreads computed under the Merton model (and using the parameters and method given in Chen et al. 2009). These spreads are very similar to, and just a little higher than, the Black-Cox spreads. What

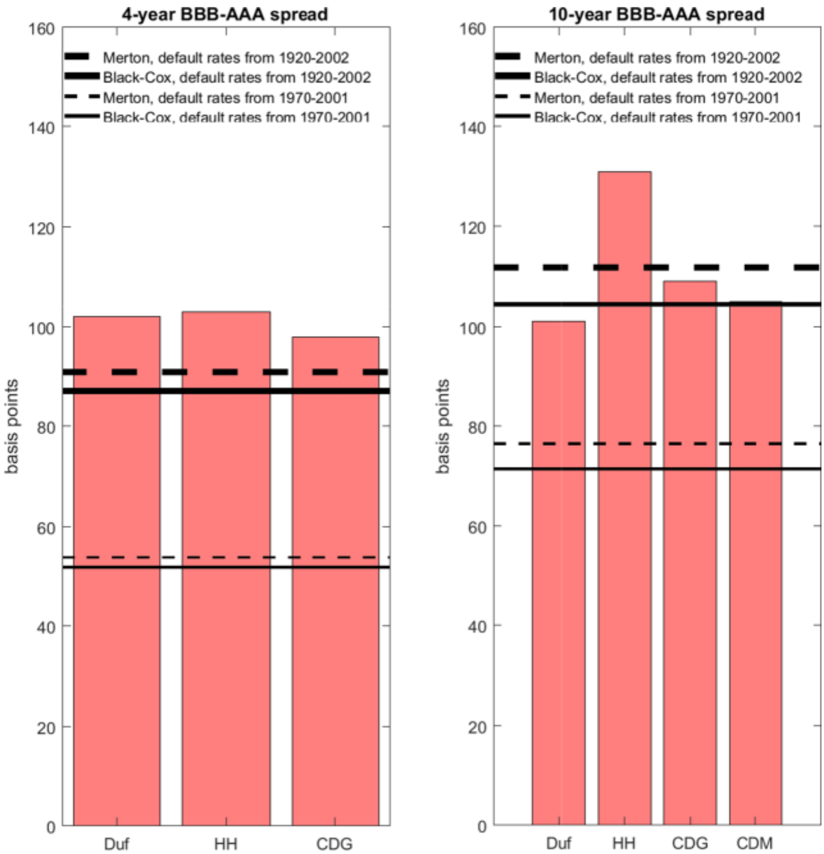

Figure 1
Actual and model-implied BBB-AAA corporate bond yield spreads when using existing approach in the literature
This figure shows actual and model-implied BBB-AAA spreads based on different estimates of actual and model-implied spreads. The actual BBB-AAA yield spreads are estimates from Duffee (1998) (Duf), Huang and Huang (2012) (HH), Chen, Collin-Dufresne, and Goldstein (2009) (CDG), and Cremers, Driessen, and Maenhout (2008) (CDM). The solid lines show spreads in the Black-Cox model based on Moody's default rates from the period 1920–2002 and 1970–2001, respectively. The dashed lines show spreads in the Merton model based on Moody's default rates from the period 1920–2002 and 1970–2001, respectively.

remains unchanged is the finding that the appearance of a credit spread puzzle depends on the sample period.

In the example we compare corporate bond yields relative to AAA yields to be consistent with CDG and others. In our later analysis we use bond yields relative to swap rates. The average difference between swap rates and AAA yields is small: over our sample period 1987–2012, the average 5- and 10-year AAA-swap spreads are 4 and 6 bps, respectively. We use swap rates in our later analysis, because the term structure of swap rates is readily available on a daily basis. There are very few AAA-rated bonds in the later part of our sample period, so we would not be able to calculate a AAA yield at different maturities.

In summary, realized average default rates vary substantially over time, and, as a result, when these are taken as ex ante default probabilities the historical period over which they are measured has a strong influence on whether or not there will appear to be a credit spread puzzle. In the next section we first explore the statistical uncertainty of historical default rates in more detail and then propose a different approach to estimating default probabilities that exploits the information contained in historical default rates more efficiently than has been the case in the literature so far.

2. Estimating Ex Ante Default Probabilities

The existing literature on the credit spread puzzle and, more broadly, the literature on credit risk typically uses the average ex post historical default rate for a single maturity and rating as an estimate of the ex ante default probability for this same maturity and rating.[4] We find that the statistical uncertainty associated with these estimates is large and propose a new approach that uses historical default rates for all maturities and ratings simultaneously to extract the ex ante default probability for any given maturity and rating. Simulations show that our approach greatly reduces statistical uncertainty.

2.1 Existing approach: Extracting the ex ante default probability from a single ex post default frequency

An ex post realized default frequency may be an unreliable estimate of the ex ante default probability for two significant reasons.

The first is that the low level of default frequency, particularly for investment-grade firms, leads to a sample size problem with default histories as short as those typically used in the literature when testing standard models. The second is that, even though the problem of sample size is potentially mitigated by the presence of a large number of firms in the cross-section, defaults are correlated across firms and so the benefit of a large cross-section in improving precision is greatly reduced.

How severe are these statistical issues? We address this question in a simulation study and base our simulation parameters on the average 10-year BBB default rate of 5.09% over 1970–2001 used in Chen, Collin-Dufresne, and Goldstein (2009). In an economy in which the ex ante 10-year default probability is 5.09% for all firms, we simulate the ex post realized 10-year default frequency over 31 years. We assume that in year 1 we have 445 identical firms, equal to the average number of firms in Moody's BBB cohorts over the

[4] Examples include Chen, Cui, He, and Milbradt (2017), Gomes, Jermann, and Schmid (2016), Christoffersen and Elkamhi (2017), Bai (2016), Zhang, Zhou, and Zhu (2009), Chen (2010), Leland (2006), Cremers, Driessen, and Maenhout (2008), Chen, Collin-Dufresne, and Goldstein (2009), Chen (2010), Huang and Huang (2012), Campello, Chen, and Zhang (2008), and McQuade (2013).

period 1970–2001. In the Black-Cox (and Merton) model firm i's asset value under the natural measure follows a GBM:

$$\frac{dV_{it}}{V_{it}} = (\mu - \delta)dt + \sigma dW_{it}^{P}, \tag{2}$$

where μ is the drift of firm value before payout of the dividend yield δ and σ is the volatility of firm value. Like in Section 1, we use our average parameter values for the period 1987–2012 estimated in Section 3: $\mu = 10.05\%$, $\delta = 4.72\%$, and $\sigma = 24.6\%$. We introduce systematic risk by assuming that

$$W_{it}^{P} = \sqrt{\rho} W_{st} + \sqrt{1-\rho} W_{it}, \tag{3}$$

where W_i is a Wiener process specific to firm i, W_s is a Wiener process common to all firms, and ρ is the pairwise correlation between percentage changes in firm value. All the Wiener processes are independent. The firm defaults the first time asset value hits a boundary equal to a fraction d of the face value of debt F, that is, the first time $V_\tau \leq dF$. The realized 10-year default frequency in the year 1 cohort is found by simulating one systematic and 445 idiosyncratic processes in Equation (3).

In year 2 we form a cohort of 445 new firms. The firms in year 2 have characteristics that are identical to those of the year 1 cohort at the time of formation. We calculate the realized 10-year default frequency of the year 2 cohort as we did for the year 1 cohort. Crucially, the common shock for years 1–9 for the year 2 cohort is the same as the common shock for years 2–10 for firms in the year 1 cohort. We repeat the same process for 22 years and calculate the overall average realized cumulative 10-year default frequency in the economy by taking an average of the default frequencies across the 22 cohorts. Finally, we repeat this entire simulation 25,000 times.

To estimate the correlation parameter ρ, we calculate pairwise equity correlations for rated industrial firms in the period 1987–2012. Specifically, for each year we calculate the average pairwise correlation of daily equity returns for all industrial firms for which Standard & Poor's provide a rating and then calculate the average of the 26 yearly estimates over 1987–2012. We estimate ρ to be 20.02%.

To set the default boundary, we proceed as follows. First, without loss of generality, we assume that the initial asset value of each firm is equal to one. This means that the firm's leverage, $L \equiv \frac{F}{V} = F$, and we set the default boundary $dF (= dL)$ such that the model-implied default probability given in Equation (A2) in the appendix matches the 10-year default rate of 5.09%.[5]

Panel A of Figure 2 shows the distribution of the realized average 10-year default rate in the simulation study and the black vertical line shows the ex

[5] We simulate firm values on a weekly basis. There is a small downward bias in default rates because a default only occurs on a weekly basis and not continuously, and we adjust for this bias by multiplying average default rates in each of the 25,000 simulations with 5.09% divided by the average of the 25,000 average 10-year BBB default rates.

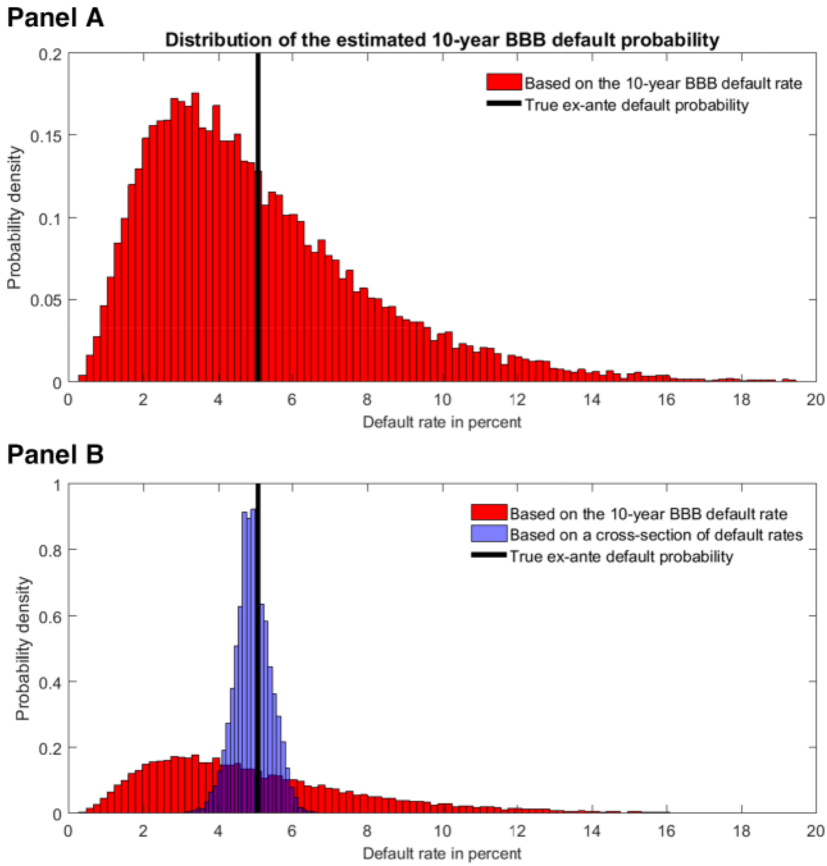

Figure 2
Distribution of estimated 10-year BBB default probability when using default rates measured over 31 years
The existing approach in the literature is to use an average historical default rate for a specific rating and maturity as an estimate for the default probability when testing spread predictions of structural models. One example is Chen, Collin-Dufresne, and Goldstein (2009), who use the 10-year BBB default rate of 5.09% realized over the period 1970-2001 as an estimate for the 10-year BBB default probability. Panel A shows the distribution of the 10-year BBB default probability when using a 31-year history of the 10-year BBB default rate as an estimate. Besides this distribution, panel B also shows the distribution of the estimated 10-year BBB default probability when extracted using the proposed approach in Section 2.2. Specifically, the default probability is estimated using the Black-Cox model and 1-, 2-, . . ., 20-year default rates for ratings AAA, ..., C averaged over 31 years.

ante default probability of 5.09%. The 95% confidence interval for the realized average default rate is wide at [1.15%; 12.78%]. We also see that the default frequency is significantly skewed to the right; that is, the modal value of around 3% is significantly below the mean of 5.09%. This means that the default frequency *most often* observed—for example, the estimate from the rating agencies—is below the mean. Specifically, although the true 10-year default probability is 5.09%, the probability that the observed average 10-year default rate over 31 years is *half* that level or less is 19.9%. This skewness means that the number reported by Moody's (5.09%) is more likely to be below the

Table 1
Distribution of average 10-year BBB default rate when measured over 31 years for different levels of systematic risk (ρ): the pairwise correlation between firms' asset values

Systematic risk		Quantiles						
ρ	Mean	0.01	0.025	0.25	0.5	0.75	0.975	0.99
0%	5.09%	4.57%	4.66%	4.94%	5.09%	5.24%	5.53%	5.62%
5%	5.09%	2.45%	2.74%	4.06%	4.93%	5.93%	8.38%	9.19%
10%	5.09%	1.69%	2.00%	3.58%	4.75%	6.22%	10.13%	11.48%
15%	5.09%	1.17%	1.49%	3.20%	4.60%	6.45%	11.41%	13.28%
20%	5.09%	0.87%	1.16%	2.86%	4.39%	6.60%	12.88%	15.20%
25%	5.09%	0.64%	0.90%	2.60%	4.23%	6.63%	14.16%	17.14%
30%	5.09%	0.44%	0.64%	2.24%	4.04%	6.80%	15.71%	18.83%
35%	5.09%	0.31%	0.48%	2.00%	3.79%	6.79%	16.88%	21.08%
40%	5.09%	0.18%	0.33%	1.73%	3.58%	6.83%	18.56%	22.82%
45%	5.09%	0.12%	0.23%	1.46%	3.32%	6.87%	19.95%	25.00%
50%	5.09%	0.08%	0.16%	1.28%	3.12%	6.86%	20.79%	25.90%

In the benchmark simulations we use an estimated value of ρ of 20.02% and simulate an average 10-year BBB default rate (based on 31 years of 10-year BBB default rates). We repeat this simulation 25,000 times and calculate the distribution of the average 10-year BBB rate. This table shows the distribution for different levels of systematic risk.

true mean than above it and, in this case, if spreads reflect the true expected default probability, they will appear too high relative to the observed historical loss rate.

There is a tradition in the literature for matching the historical 10-year BBB default rate exactly by backing out one or several model parameters (see footnote 4 for references). In this tradition, the model-implied default probability will inherit one-to-one the statistical uncertainty of the historical default rate. Since the statistical uncertainty of the historical default rate is large, the statistical uncertainty of other important model predictions, such as the predicted spread, will likewise be large.

Given that we simulate 9,345 firms over a period of 31 years, it might be surprising that the realized default rate can be far from the ex ante default probability. The reason is simply the presence of systematic risk in the economy which induces correlation in defaults across firms. Table 1 shows the results of the simulation described above for different levels of systematic risk (ρ). With zero systematic risk (first row of the table) the distribution of the 10-year default rate is naturally tightly centered on the true mean with a 95% confidence interval of [4.66%, 5.53%]. However, the dispersion in realized default rates increases substantially, even with modest levels of systematic risk. When ρ is 10%, for example, half our estimated value of 20.02%, the 95% confidence interval becomes [2.00%, 10.13%], which is 70% as wide as when $\rho = 20.02\%$. For higher values of ρ, the confidence interval is wider still.

Moody's now publish default rates starting from 1920 and, other things equal, a longer time period over which average default rates are measured will lead to improved statistical precision. However, even if the average default rate is measured over 92 years there remains significant statistical uncertainty when estimates of the default probability are based on a single rating and maturity. Keeping the default probability fixed at 5.09% (and with $\rho = 20.02\%$)

and increasing the simulated time period from 31 to 92 years, leads to a 95% confidence interval of [2.47%; 8.95%]. Thus, even with 92 years of default data, there is still significant uncertainty regarding the true ex ante default probability.

2.2 A new approach: Extracting the ex ante default probability from a cross-section of ex post default frequencies

We now describe a method for estimating default probabilities that uses realized cumulative default rates from a wide range of ratings and maturities. The key feature of our method that allows us to aggregate default rate information across ratings and maturities is the assumption that firms with different ratings and having bonds with different maturities will nonetheless share a common default boundary.

The statistical benefits of the new approach derive from two main sources. First, because low credit quality bonds default more frequently, they provide much more information on the location of the default boundary than high credit quality bonds, and we can therefore obtain better estimates of the default probability on the latter when we also include default rate information on the former. Second, estimates of the default boundary obtained from different maturity-rating pairs are imperfectly correlated, and it is therefore efficient to combine them. Both these effects lead to better precision in estimated default probabilities than is obtained from a single default rate.

Our approach is a generalization of the method used by Chen, Collin-Dufresne, and Goldstein (2009), who find the default boundary parameter d such that the model-implied BBB default probability matches the average historical 10-year BBB default rate. Given estimates of $\Theta^P = (\mu, \sigma, \delta)$ and the leverage for BBB-rated firms, L_{BBB}, this approach amounts to finding d such that the model-implied default probability $\pi^P(dL_{BBB}, \Theta^P, 10)$ (given in Equation (A2) in the appendix) is equal to the historical 10-year BBB default rate. In other words, they find d as

$$\min_{\{d\}} \left| \pi^P(dL_{BBB}, \Theta^P, 10) - \hat{\pi}^P_{BBB,10} \right|, \tag{4}$$

and since they only match one historical default rate, the error in the objective function in expression (4) is zero. Chen, Collin-Dufresne, and Goldstein (2009) then use this default boundary to calculate the 10-year BBB spread according to Equation (1).

Our approach is similar but, crucially, we fit to the historical default rates on *all* available ratings and maturities. We estimate the default boundary parameter d by minimising the sum of absolute deviations between annualized model-implied and historical default rates:

$$\min_{\{d\}} \sum_{a=AAA}^{C} \sum_{T=1}^{20} \frac{1}{T} \left| \pi^P(dL_a, \Theta^P, T) - \hat{\pi}^P_{aT} \right|, \tag{5}$$

where $\hat{\pi}^P_{aT}$ is the historical cumulative default rate for rating a and maturity T. Since we fit to a range of historical default rates, each individual default rate is fitted with error.[6] If we restrict the objective function in (5) to include only the 10-year BBB default rate, then our approach is identical to that of Chen, Collin-Dufresne, and Goldstein (2009).

We assume in expression (5) that the default boundary parameter d is homogeneous across firms. This is a standard assumption in the literature and the majority of the models considered in Eom, Helwege, and Huang (2004) and Huang and Huang (2012) are estimated under this assumption. Structural models with endogenous default such as the Leland (1994) model have a default boundary that depends on a number of parameters and it would be straightforward to allow for this in the objective function given in Equation (5).

How do the results from using the proposed criterion in expression (5) compare with those from using (4)? We answer this question by simulating cumulative default rates and then, for each drawing, estimating d and computing the 10-year BBB default probability.

Using the objective function given in (4) the realized 10-year default rate is fitted without error and so in this case the distribution of the estimated ex ante 10-year BBB default probability is simply the distribution of the *realized* default rate given in panel A of Figure 2.

In the case of the objective function given in (5) we proceed as follows. Following the procedure described in the previous section, we simulate over a period of 31 years but now for each of the seven major rating categories. We assume that all firms, regardless of rating, have the same parameters σ, δ, and μ, the same correlation structure (given in Equation (3)) and the same correlation parameter, ρ.

We assume the default boundary to be $d=1$ and the leverage for each rating category is set, like in the previous section, such that the historical 10-year cumulative default rate is matched.[7] Since d is common across all firms, it means that firms in different rating categories differ only in their initial leverage. We then simulate 1-, . . ., 20-year default rates. We set the number of firms in each rating cohort equal to the average number of firms in the Moody's cohorts for that rating category in the period 1970–2001. We carry out this simulation 25,000 times and for each simulation we first estimate d according to expression (5) and then calculate the 10-year BBB cumulative default probability in Equation (A2).

Panel B of Figure 2 shows the distribution of the 10-year BBB estimated default probability under the objective function given in Equation (5) and, for comparison, the distribution obtained using the criterion in expression (4) and already given in panel A. We emphasize that, although each estimate of d in the first case (panel A) results in an estimate of the 10-year BBB default probability

[6] There may be weighting patterns other than $\frac{1}{T}$ that are even more efficient, but, in making this choice, our objective has been to keep our method as simple and transparent as possible.

[7] Since the default boundary d and leverage L only enter the default probability as a product, any value of d gives rise to the same simulation results because leverage adjusts.

that matches the realized default rate exactly, the default probabilities in both distributions are consistent with the Black-Cox formula (A2), and the difference in the distributions solely arises from the application of either the fitting criterion (4) (panel A) or the new method (5) (panel B) in estimating the default boundary from the *same set* of realized default rates.

We see that the distribution is both much tighter and less skewed than when the estimate is solely based on the 10-year BBB default rate. As we show next, the distribution is tighter because we use default rates from all maturities and ratings, in particular from low ratings, instead of just one. The skewness is reduced because we include default rates in the estimation that are significantly higher than BBB default rates. To see why this is the case, consider the example discussed earlier of a large number of firms with a default probability of 5% and where their defaults are perfectly correlated. In this case we will see no defaults 95% of the time and the realized default rate will underestimate the default probability 95% of the time and the distribution will exhibit positive skewness. By the same logic, a higher default probability reduces the skewness and for default probabilities greater than 50% the skewness is in fact negative.

Table 2 provides summary statistics—standard deviation, skewness, and quantiles at 2.5%, 50%, and 97.5%—on the distribution of the estimated 10-year BBB default probability derived from the simulations described above. Results are given for a number of cases in which different subsets of ratings and maturities are used in the estimation.

Panel A of Table 2 gives results for the value of systematic risk estimated from the data ($\rho = 0.2002$). The first row in this panel gives results for the case that includes all ratings and maturities and corresponds to our implementation of this method using actual firm and default rate data (described in the next section). The second row corresponds to the previous literature that estimates the default boundary from a single average 10-year BBB default rate. The difference between the two distributions is striking: using all ratings and maturities leads to a standard deviation that is around 84% lower and has almost no skewness (compared with a skewness of 1.28 using a single BBB default rate).

What is responsible for this improvement? By searching for a parameter, d, that is common to all ratings our approach allows us to aggregate information across ratings and maturities. As suggested earlier, the inclusion of default data on high-yield debt is particularly productive and row 3 shows the benefit from including the 10-year default rate from each of the seven ratings rather than just one. Relative to estimates derived from the 10-year BBB default rate alone (row 2), this step alone almost halves the standard deviation and substantially reduces the skewness. Including all (105) ratings for maturities longer than 5 years (row 4) further reduces both the standard deviation and skewness but including only default rates for longer maturities (longer than 10 years, row 5) is counterproductive. Further simulation experiments suggest that this is because long-maturity default rates for highly rated debt have both high dispersion and high skewness.

Table 2
Properties of the 10-year BBB default probability estimator when using a subset of realized default rates in the estimation

	SD	Skewness	Quantiles 0.025	0.5	0.975
A. Systematic risk $\rho = 0.2002$					
All ratings, all maturities (140 default rates)	**0.48**	**−0.05**	**3.96%**	**4.92%**	**5.88%**
10-year BBB (1 default rate)	3.05	1.28	1.15%	4.40%	12.78%
All ratings, 10-year (7 default rates)	1.65	0.55	2.45%	5.03%	8.77%
All ratings, maturity >5 years (105 default rates)	1.22	0.43	3.02%	5.03%	7.70%
All ratings, maturity >10 years (70 default rates)	2.01	0.82	2.15%	5.06%	10.07%
Investment grade, all maturities (80 default rates)	3.77	1.21	0.45%	4.33%	14.70%
Investment grade, maturity ≤10 years (40 default rates)	2.39	0.95	1.45%	4.62%	10.63%
Increased number of 10-year BBB observations	2.66	1.17	1.41%	4.53%	11.69%
B. Systematic risk $\rho = 0.1$					
All ratings, all maturities (140 default rates)	**0.35**	**0.16**	**4.46%**	**5.07%**	**5.79%**
10-year BBB (1 default rate)	2.07	1.03	2.04%	4.74%	9.97%
All ratings, 10-year (7 default rates)	1.22	0.45	3.33%	5.36%	8.12%
All ratings, maturity >5 years (105 default rates)	0.92	0.31	3.76%	5.35%	7.31%
All ratings, maturity >10 years (70 default rates)	1.49	0.67	3.13%	5.46%	8.98%
Investment grade, all maturities (80 default rates)	2.74	0.92	1.37%	4.95%	11.93%
Investment grade, maturity ≤10 years (40 default rates)	1.67	0.68	2.45%	4.95%	8.98%
Increased number of 10-year BBB observations	1.76	0.82	2.36%	4.84%	9.19%
C. Systematic risk $\rho = 0.3$					
All ratings, all maturities (140 default rates)	**0.59**	**−0.01**	**3.90%**	**5.06%**	**6.26%**
10-year BBB (1 default rate)	3.98	1.77	0.66%	4.00%	15.51%
All ratings, 10-year (7 default rates)	2.16	0.70	2.15%	5.31%	10.52%
All ratings, maturity >5 years (105 default rates)	1.60	0.51	2.79%	5.30%	8.77%
All ratings, maturity >10 years (70 default rates)	2.72	0.99	1.87%	5.41%	12.41%
Investment grade, all maturities (80 default rates)	4.72	1.47	0.16%	3.82%	17.44%
Investment grade, maturity ≤10 years (40 default rates)	3.07	1.25	0.80%	4.24%	12.41%
Increased number of 10-year BBB observations	3.40	1.55	0.98%	4.28%	13.65%

This table shows the properties of the estimated 10-year BBB default probability (as a percentage) when the estimate is based on average historical default rates over 31 years. The table is based on 25,000 simulations of average default rates over 31 years. ‘All ratings, all maturities’ refers to the method proposed in Section 2.2, where the default probability estimate for a single rating and maturity is extracted using the Black-Cox model and 1-, 2-, ..., 20-year default rates for ratings AAA, AA, A, BBB, BB, B, and C. ‘10-year BBB’ refers to the standard approach in the literature of using the historical average 10-year BBB default rate as an estimate of the ex ante 10-year BBB default probability. ’increased number of 10-year BBB observations’ refers to using only the 10-year BBB default rate as an estimator of the 10-year default probability, but increasing the number of BBB firms in each cohort to 40,660 instead of 445 like in the benchmark case. In the remaining cases some, but not all, default rates are used when estimating the 10-year BBB default probability.

These characteristics of long-maturity default rates for high ratings emerge strongly if we restrict the default data used to estimate d to investment-grade debt. Row 6 shows that using investment-grade default rates for all maturities actually leads to lower precision than using only the 10-year BBB default rate, while row 7 shows that this problem disappears when we include investment-grade default rates for maturities of up to 10 years but exclude those for longer maturities.

In the simulations the number of firms in each rating bucket is the average number in Moody’s cohort over the period 1970–2001. The distribution of the number of firms across ratings and over time will also influence the results but this effect does not appear to be very strong. The average number of firms in the BBB bucket, for example, is 445. Row 8 in panel A of Table 2 asks the question of how much the estimate of the default rate would be improved if the number of firms in the BBB rating bucket each year were equal to the 40,660 (the average number of firms in Moody's cohorts across all ratings in any 1 year—2033—multiplied by 20, the number of horizons). As the results show, increasing the number of firms in the BBB cohort by over 90 times results in only a modest reduction in the standard deviation: 2.66 versus 3.08.

Panels B and C of Table 2 report the corresponding results for values of ρ of 0.1 and 0.3. The results have much the same character in both cases. When $\rho = 0.1$ the standard deviation in each case is about 30% lower and with $\rho = 0.3$ about 30% higher. The rank order of the eight results is the same for all three values of ρ.

The results in Table 2 show that, compared with previous methods, the improvement in the precision of the estimated default probability comes from the use of default rates from a *range* of ratings and maturities and, specifically, the use of default rates on high-yield debt. As mentioned earlier, the results also suggest that refinement of the weighting scheme in Equation (5) might lead to even lower standard errors. However, the assumption of equal weights in our approach scheme is both simple and transparent (and, quite possibly, more robust than weights that are "optimized" in the context of a particular model).

Tables 3 and 4 give further details on the performance of our method. Table 3 compares the standard deviation of the estimated default probabilities in the existing approach and our approach for all the ratings and maturities we study in the next section. Our approach results in a substantial reduction in the standard deviation, except for the shortest maturities of the lowest ratings. Table 4 shows that for investment-grade ratings, our approach also results in a substantial reduction in the skewness.

Overall, our proposed method greatly reduces both the standard deviation and skewness of estimated investment-grade default probabilities.[8]

3. A New Perspective on the Credit Spread Puzzle

In the previous section we proposed a new method to estimate default probabilities by combining information on historical default rates from the cross-section of rating categories and maturities and showed that, compared to the existing approach of using a single default rate, it greatly improves statistical precision. In this section we apply our method to a large data set of bond quotes over 1987–2012 to shed new light on the credit spread puzzle.

[8] In our simulations we match Moody's historical 10-year default rates over 1970–2001 for ratings AAA, . . ., C. In the Internet Appendix we show that the distribution of the default boundary parameter d—and therefore also the distribution of the default probability—is similar if we match Moody's default rates at maturities other than 10 years.

Table 3
Standard deviation of the estimated default probability

Horizon (years)	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
AAA																	
New method	0.00	0.01	0.02	0.03	0.05	0.07	0.09	0.11	0.14	0.16	0.18	0.20	0.22	0.24	0.26	0.28	0.29
Existing method	0.02	0.05	0.11	0.21	0.34	0.52	0.72	0.95	1.20	1.46	1.73	2.00	2.27	2.54	2.79	3.03	3.27
Ratio	10%	14%	15%	15%	14%	13%	13%	12%	11%	11%	10%	10%	10%	10%	9%	9%	9%
AA																	
New method	0.00	0.01	0.02	0.04	0.06	0.08	0.10	0.13	0.15	0.17	0.20	0.22	0.24	0.26	0.28	0.29	0.31
Existing method	0.01	0.05	0.11	0.22	0.37	0.56	0.78	1.03	1.29	1.56	1.85	2.13	2.41	2.68	2.94	3.19	3.44
Ratio	17%	18%	17%	16%	15%	14%	13%	12%	12%	11%	11%	10%	10%	10%	9%	9%	9%
A																	
New method	0.01	0.02	0.04	0.07	0.10	0.14	0.17	0.20	0.23	0.26	0.29	0.32	0.34	0.36	0.38	0.40	0.42
Existing method	0.03	0.10	0.22	0.40	0.64	0.91	1.23	1.57	1.91	2.26	2.62	2.97	3.30	3.63	3.94	4.23	4.51
Ratio	24%	22%	20%	18%	16%	15%	14%	13%	12%	12%	11%	11%	10%	10%	10%	9%	9%
BBB																	
New method	0.08	0.15	0.22	0.29	0.36	0.42	0.48	0.52	0.57	0.60	0.64	0.67	0.69	0.71	0.73	0.75	0.77
Existing method	0.24	0.53	0.93	1.40	1.92	2.47	3.05	3.63	4.19	4.72	5.25	5.76	6.22	6.66	7.06	7.44	7.78
Ratio	32%	27%	24%	21%	19%	17%	16%	14%	14%	13%	12%	12%	11%	11%	10%	10%	10%
BB																	
New method	0.98	1.17	1.31	1.40	1.47	1.52	1.55	1.58	1.59	1.60	1.61	1.61	1.62	1.62	1.61	1.61	1.61
Existing method	2.27	3.28	4.23	5.15	6.02	6.84	7.64	8.38	9.09	9.76	10.40	10.99	11.55	12.06	12.53	12.94	13.32
Ratio	43%	36%	31%	27%	24%	22%	20%	19%	18%	16%	15%	15%	14%	13%	13%	12%	12%
B																	
New method	3.11	3.09	3.05	2.99	2.93	2.87	2.82	2.77	2.72	2.68	2.64	2.60	2.56	2.53	2.50	2.47	2.44
Existing method	5.35	6.27	7.06	7.77	8.41	9.01	9.55	10.06	10.55	11.02	11.47	11.89	12.28	12.65	12.98	13.30	13.59
Ratio	58%	49%	43%	38%	35%	32%	30%	28%	26%	24%	23%	22%	21%	20%	19%	19%	18%
C																	
New method	5.21	4.84	4.56	4.33	4.15	4.00	3.87	3.75	3.65	3.57	3.49	3.42	3.35	3.30	3.24	3.19	3.15
Existing method	5.15	5.45	5.72	5.98	6.23	6.45	6.67	6.87	7.09	7.29	7.49	7.68	7.86	8.03	8.19	8.36	8.50
Ratio	101%	89%	80%	72%	67%	62%	58%	55%	52%	49%	47%	44%	43%	41%	40%	38%	37%

This table shows the standard deviation of the estimated default probability (as a percentage) when the estimate is based on average historical default rates over 31 years. The table is based on 25,000 simulations of average default rates over 31 years. 'Existing method' refers to the standard approach in the literature of using the average historical default rate for a single rating and maturity as an estimate of the ex ante default probability. 'New method' refers to the method proposed in Section 2.2 where the default probability estimate for a single rating and maturity is extracted using the Black-Cox model and 1-, 2-, ..., 20-year default rates for ratings AAA, AA, A, BBB, BB, B, and C. 'Ratio' is the standard deviation of default probability estimates using the new method divided by the standard deviation of default probability estimates using the existing method.

Table 4
Skewness of the estimated default probability

Horizon (years)	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
AAA																	
New method	0.51	0.36	0.25	0.18	0.12	0.07	0.04	0.01	−0.02	−0.04	−0.06	−0.07	−0.09	−0.10	−0.11	−0.12	−0.13
Existing method	3.57	2.87	2.56	2.35	2.23	2.05	1.90	1.76	1.63	1.51	1.40	1.30	1.21	1.13	1.05	1.00	0.94
Ratio	14%	12%	10%	8%	5%	4%	2%	0%	−1%	−3%	−4%	−6%	−7%	−9%	−10%	−12%	−14%
AA																	
New method	0.50	0.35	0.24	0.17	0.11	0.07	0.03	0.00	−0.02	−0.04	−0.06	−0.08	−0.09	−0.10	−0.11	−0.12	−0.13
Existing method	3.35	3.04	2.69	2.44	2.23	2.06	1.91	1.77	1.63	1.51	1.39	1.29	1.21	1.12	1.05	1.00	0.94
Ratio	15%	11%	9%	7%	5%	3%	2%	0%	−1%	−3%	−4%	−6%	−7%	−9%	−11%	−12%	−14%
A																	
New method	0.44	0.30	0.20	0.13	0.08	0.04	0.01	−0.02	−0.04	−0.06	−0.08	−0.09	−0.11	−0.12	−0.13	−0.14	−0.14
Existing method	2.89	2.56	2.36	2.17	2.00	1.87	1.73	1.60	1.49	1.38	1.28	1.19	1.11	1.04	0.98	0.92	0.87
Ratio	15%	12%	9%	6%	4%	2%	0%	−1%	−3%	−5%	−6%	−8%	−10%	−11%	−13%	−15%	−17%
BBB																	
New method	0.29	0.18	0.10	0.05	0.01	−0.03	−0.05	−0.08	−0.09	−0.11	−0.12	−0.13	−0.14	−0.15	−0.16	−0.17	−0.17
Existing method	1.75	1.63	1.54	1.48	1.40	1.34	1.28	1.21	1.15	1.09	1.02	0.95	0.90	0.84	0.79	0.74	0.69
Ratio	17%	11%	7%	3%	0%	−2%	−4%	−6%	−8%	−10%	−12%	−14%	−16%	−18%	−20%	−23%	−25%
BB																	
New method	0.08	0.00	−0.05	−0.08	−0.11	−0.13	−0.14	−0.16	−0.17	−0.18	−0.19	−0.19	−0.20	−0.20	−0.21	−0.21	−0.22
Existing method	0.73	0.69	0.66	0.64	0.63	0.63	0.62	0.61	0.60	0.58	0.57	0.55	0.53	0.50	0.47	0.44	0.41
Ratio	10%	0%	−7%	−13%	−17%	−20%	−23%	−26%	−28%	−30%	−33%	−35%	−38%	−41%	−44%	−48%	−53%
B																	
New method	−0.08	−0.12	−0.15	−0.17	−0.19	−0.20	−0.21	−0.21	−0.22	−0.22	−0.23	−0.23	−0.24	−0.24	−0.24	−0.24	−0.24
Existing method	0.24	0.21	0.19	0.17	0.16	0.15	0.14	0.13	0.12	0.10	0.09	0.08	0.07	0.06	0.04	0.02	−0.00
Ratio	−35%	−60%	−79%	−99%	−115%	−132%	−149%	−170%	−189%	−215%	−241%	−278%	−330%	−404%	−605%	−1301%	28766%
C																	
New method	−0.20	−0.21	−0.23	−0.23	−0.24	−0.24	−0.25	−0.25	−0.25	−0.25	−0.26	−0.26	−0.26	−0.26	−0.26	−0.26	−0.26
Existing method	−0.10	−0.13	−0.16	−0.18	−0.21	−0.22	−0.24	−0.26	−0.28	−0.30	−0.32	−0.34	−0.36	−0.38	−0.41	−0.44	−0.47
Ratio	195%	159%	142%	129%	116%	109%	102%	95%	89%	85%	80%	76%	72%	68%	64%	59%	56%

This table shows the skewness of the estimated default probability (as a percentage) when the estimate is based on average historical default rates over 31 years. The table is based on 25,000 simulations of average default rates over 31 years. ‘Existing method’ refers to the standard approach in the literature of using the average historical default rate for a single rating and maturity as an estimate of the ex ante default probability. ‘New method’ refers to the method proposed in Section 2.2, where the default probability estimate for a single rating and maturity is extracted using the Black-Cox model and 1-, 2-, ..., 20-year default rates for ratings AAA, AA, A, BBB, BB, B, and C. ‘Ratio’ is the skewness of default probability estimates using the new method divided by the skewness of default probability estimates using the existing method.

3.1 Data

For the period January 1, 1997 to July 1, 2012, we use daily quotes provided by Merrill Lynch (ML) on all corporate bonds included in the ML investment-grade and high-yield indices. These data are used by Schaefer and Strebulaev (2008) and Acharya, Amihud, and Bharath (2013), among others. We obtain bond information from the Mergent Fixed Income Securities Database (FISD) and limit the sample to senior unsecured fixed rate or zero coupon bonds. We exclude bonds that are callable, are convertible, are putable, are perpetual, are foreign denominated, are Yankee, have sinking fund provisions, or have covenants.[9] For the period April 1987 to December 1996 we use monthly data from the Lehman Brothers Fixed Income Database. These data are used by Duffee (1998), Huang and Huang (2012), and Acharya, Amihud, and Bharath (2013), among others. We include only data from the Lehman database that are actual quotes (in contrast to data based on matrix-pricing). The Lehman database starts in 1973, but there are two reasons we start from April 1987. First, there are few noncallable bonds before the mid-1980s (see Duffee 1998) and second, we calculate credit spreads relative to the swap rate and we do not have data on swap rates prior to April 1987. We use only bonds issued by industrial firms and restrict our sample to bonds with a maturity of less than 20 years to be consistent with the maturities of the default rates we use as part of the estimation.[10] In total we have 256,698 observations. We show in the Internet Appendix that dealer quotes are unreliable for short-maturity bonds due to quotes being bid quotes and when we report bond spreads we therefore exclude bonds with a maturity less than 3 years.

Table 5 shows summary statistics for the corporate bond sample. The table shows that the number of bonds with a low rating of B or C is small; for example there is only one C-rated bond in the maturity group 13–20 years. The reason is that speculative-grade bonds frequently contain call options which leads to their exclusion from our sample (see also Booth et al. 2014). In the following we report results for ratings B and C with the caveat that these results—particularly for long maturities—are based on few observations and therefore are noisy.

To price a bond in the Black-Cox model, we need the issuing firm's asset volatility, leverage ratio, and payout ratio along with the bond's recovery rate. *Leverage ratio* is calculated as the book value of debt divided by firm value (where firm value is calculated as book value of debt plus market value of equity). *Payout ratio* is calculated as the sum of interest payments to debt, dividend payments to equity, and net stock repurchases divided by firm value.

[9] For bond rating, we use the lower of Moody's rating and S&P's rating. If only one of the two rating agencies have rated the bond, we use that rating. We track rating changes on a bond, so the same bond can appear in several rating categories over time.

[10] In the Internet Appendix we show that our results are similar if we use TRACE transactions data for the shorter period of 2002–2012.

Table 5
Bond summary statistics

3- to 7-year bond maturity	AAA	AA	A	BBB	BB	B	C	all
Number of bonds	21	109	327	289	116	36	14	753
Mean number of bonds pr month	1.69	7.54	22.3	18	6.94	1.71	0.37	58.4
Mean number of quotes pr month	1.687	45.68	105.7	109.8	57.99	13.31	4.497	338.7
Age	2.05	4.73	5.19	5.89	4.37	3.07	2.50	5.08
Coupon	7.19	6.30	7.00	7.57	7.62	9.61	11.72	7.36
Amount outstanding ($mm)	265	329	277	321	414	300	231	322
Time-to-maturity	4.74	4.73	4.87	4.77	4.92	5.25	5.23	4.85
7- to 13-year bond maturity	AAA	AA	A	BBB	BB	B	C	all
Number of bonds	16	94	288	276	100	26	11	680
Mean number of bonds pr month	1.28	7.98	20.9	18.1	5.04	0.93	0.3	54.4
Mean number of quotes pr month	1.913	28.19	67.3	70.52	27.12	5.61	3.35	204
Age	7.19	6.95	5.67	4.88	2.86	5.21	9.46	5.26
Coupon	7.21	7.29	7.30	7.84	7.76	9.23	9.95	7.64
Amount outstanding ($mm)	400	308	265	294	535	222	270	317
Time-to-maturity	10.37	8.96	9.26	9.14	8.74	8.94	10.14	9.12
13- to 20-year bond maturity	AAA	AA	A	BBB	BB	B	C	all
Number of bonds	3	21	81	75	30	9	1	173
Mean number of bonds pr month	0.307	2.06	8.28	6.26	2.09	0.473	0.0233	19.5
Mean number of quotes pr month	3.73	2.06	26.04	17.94	5.447	2.243	0.41	57.87
Age	14.63	2.40	7.74	5.64	9.64	11.01	9.31	7.66
Coupon	8.29	8.04	8.17	8.25	8.41	7.28	9.22	8.19
Amount outstanding ($mm)	687	208	332	218	227	157	183	297
Time-to-maturity	15.19	16.77	16.61	15.84	16.95	17.27	14.07	16.33

The sample consists of noncallable bonds with fixed coupons issued by industrial firms. This table shows summary statistics for the data set. Bond yield quotes cover the period 1987Q2–2012Q2. 'Number of bonds' is the number of bonds that appear (in a particular rating and maturity range) at some point in the sample period. 'Mean number of bonds pr month' is the average number of bonds that appear in a month. 'Mean number of quotes pr month' is the total number of quotes in the sample period divided by the number of months. For each quote we calculate the bond's time since issuance and 'Age' is the average time since issuance across all quotes. 'Coupon' is the average bond coupon across all quotes. 'Amount outstanding' is the average outstanding amount of a bond issue across all quotes. 'Time-to-maturity' is the average time until the bond matures across all quotes.

An important parameter is the *Asset volatility*, and here we follow the approach of Schaefer and Strebulaev (2008) in calculating asset volatility. Since firm value is the sum of the debt and equity values, asset volatility is given by

$$\sigma_t^2 = (1-L_t)^2\sigma_{E,t}^2 + L_t^2\sigma_{D,t}^2 + 2L_t(1-L_t)\sigma_{ED,t}, \tag{6}$$

where σ_t is the volatility of assets, $\sigma_{D,t}$ volatility of debt, $\sigma_{ED,t}$ the covariance between the returns on debt and equity, and L_t is leverage ratio. If we assume that debt volatility is zero, asset volatility reduces to $\sigma_t = (1-L_t)\sigma_{E,t}$. This is a lower bound on asset volatility. Schaefer and Strebulaev (2008) (SS) compute this lower bound along with an estimate of asset volatility that implements Equation (6) in full. They find that for investment-grade companies the two estimates of asset volatility are similar while for junk bonds there is a significant difference. We compute the lower bound of asset volatility, $(1-L_t)\sigma_{E,t}$, and multiply this lower bound with SS's estimate of the ratio of asset volatility computed from Equation (6) to the lower bound. Specifically, we estimate $(1-L_t)\sigma_{E,t}$ and multiply this by 1 if $L_t < 0.25$, 1.05 if $0.25 < L_t \leq 0.35$, 1.10 if $0.35 < L_t \leq 0.45$, 1.20 if $0.45 < L_t \leq 0.55$, 1.40 if $0.55 < L_t \leq 0.75$, and 1.80 if $L_t > 0.75$.[11] This method has the advantage of being transparent and easy to replicate. For a given firm we then compute the average asset volatility over the sample period and use this constant asset volatility for every day in the sample period. All firm variables are obtained from CRSP and Compustat, and Appendix B provides the details.

Summary statistics for the firms in our sample are shown in Table 6. The average leverage ratios of 0.13 for AAA, 0.14 for AA, 0.27 for A, and 0.37 for BBB are similar to those found in other papers: Huang and Huang (2012) use a leverage ratio of 0.13 for AAA, 0.21 for AA, 0.32 for A, and 0.43 for BBB, and Schaefer and Strebulaev (2008) find an average leverage of 0.10 for AAA, 0.21 for AA, 0.32 for A, and 0.37 for BBB. Average equity volatility is monotonically increasing with rating, consistent with a leverage effect. The estimates are similar to those in SS for A-AAA ratings, while the average equity volatility for BBB firms of 0.38 is higher than the value of 0.33 given in SS. Asset volatilities are slightly increasing in rating and broadly consistent with the estimates in SS.

We set the recovery rate to 37.8%, which is Moody's (2013) average recovery rate, as measured by post-default trading prices, for senior unsecured bonds for the period 1982–2012. Finally, the risk-free rate, r, is the swap rate for the same maturity as the bond. Traditionally, Treasury yields have been used as risk-free rates, but recent evidence shows that swap rates are a better proxy than Treasury yields. A major reason is that Treasury bonds enjoy a convenience yield that pushes their yields below risk-free rates (Feldhütter and Lando 2008; Krishnamurthy and Vissing-Jorgensen 2012; Nagel 2016). The convenience yield is for example due to the ability to post Treasuries as collateral with a significantly lower haircut than other financial securities, an effect outside the scope of the model.

3.2 Estimation of the default boundary

In this section we estimate the default boundary, a single parameter that we then use to price bonds across ratings and maturities. Although estimating the default boundary parameter, d, by fitting to historical (natural) default rates requires an estimate of the Sharpe ratio, when we compute yield spreads we use standard "risk-neutral" pricing.

We follow the method outlined in Section 2.2. Specifically, if we observe a spread on bond i with a time-to-maturity T issued by firm j on date t, we calculate the firm's T-year default probability $\pi^P(dL_{jt}, \Theta^P_{jt}, T)$ where $\Theta^P_{jt} =$

11 These fractions are based on those in table 7 of SS apart from 1.80, which we deem to be reasonable. Results are insensitive to other reasonable choices of values for $L > 0.75$. See also Correia, Kang, and Richardson (2014) for an assessment of different approaches to calculating asset volatility.

Table 6
Firm summary statistics, industrial firms in Compustat with straight bullet bonds outstanding

	#Firms	Mean	10th	25th	Median	75th	90th
			Leverage ratio				
AAA	10	0.13	0.07	0.08	0.09	0.14	0.23
AA	53	0.14	0.07	0.10	0.13	0.17	0.23
A	170	0.27	0.13	0.17	0.24	0.34	0.46
BBB	197	0.37	0.17	0.25	0.36	0.48	0.58
BB	100	0.46	0.18	0.29	0.47	0.61	0.73
B	40	0.52	0.21	0.32	0.48	0.73	0.83
C	6	0.76	0.60	0.70	0.80	0.92	0.96
All	393	0.33	0.11	0.18	0.29	0.45	0.61
			Equity volatility				
AAA	10	0.19	0.15	0.17	0.18	0.20	0.24
AA	53	0.27	0.18	0.22	0.25	0.33	0.37
A	170	0.30	0.20	0.24	0.30	0.36	0.41
BBB	197	0.37	0.24	0.27	0.34	0.42	0.55
BB	100	0.46	0.25	0.31	0.42	0.53	0.74
B	40	0.51	0.31	0.36	0.48	0.65	0.78
C	6	0.73	0.34	0.64	0.73	0.79	1.02
All	393	0.35	0.21	0.25	0.32	0.40	0.54
			Asset volatility				
AAA	10	0.18	0.15	0.15	0.19	0.19	0.19
AA	53	0.23	0.19	0.22	0.23	0.25	0.28
A	170	0.24	0.17	0.22	0.24	0.28	0.29
BBB	197	0.25	0.17	0.19	0.24	0.28	0.36
BB	100	0.27	0.17	0.23	0.25	0.28	0.40
B	40	0.28	0.16	0.22	0.25	0.39	0.42
C	6	0.26	0.17	0.18	0.22	0.26	0.42
All	393	0.25	0.17	0.22	0.24	0.28	0.33
			Payout ratio				
AAA	10	0.036	0.012	0.017	0.030	0.049	0.071
AA	53	0.041	0.015	0.027	0.040	0.053	0.066
A	170	0.047	0.019	0.030	0.043	0.058	0.079
BBB	197	0.050	0.017	0.027	0.045	0.065	0.098
BB	100	0.045	0.020	0.027	0.040	0.057	0.078
B	40	0.046	0.018	0.029	0.041	0.061	0.077
C	6	0.068	0.037	0.049	0.056	0.096	0.103
All	393	0.047	0.018	0.028	0.043	0.059	0.083

For each bond yield observation, the leverage ratio, equity volatility, asset volatility, and payout ratio are calculated for the issuing firm on the day of the observation. Leverage ratio is the ratio of the book value of debt to the market value of equity plus the book value of debt. Equity volatility is the annualized volatility of daily equity returns from the last 3 years. Asset volatility is the unlevered equity volatility, calculated as explained in the text. Payout ratio is yearly interest payments plus dividends plus share repurchases divided by firm value. Firm variables are computed using data from CRSP and Compustat.

$(\mu_{jt}, \sigma_j, \delta_{jt})$ using Equation (A2). Here, σ_j is the firm's constant asset volatility, L_{jt}, μ_{jt} and δ_{jt} are the time-t estimates of the firm's leverage ratio, asset value drift and payout rate. To calculate μ_{jt} we assume a constant Sharpe ratio θ such that $\mu_{jt} = \theta\sigma_j + r_t^T - \delta_{jt}$, where r_t^T is the T-year risk-free rate. We use Chen, Collin-Dufresne, and Goldstein's (2009) estimate of the Sharpe ratio of 0.22. In the simulations described above we assumed that firms in a given rating category had identical initial leverage, payout ratio and asset volatility. Here, however, leverage, payout ratio and asset volatility are all firm specific.

For a given rating a and maturity T rounded up to the nearest integer year, we find all bond observations in the sample with the corresponding rating and maturity. For a given calendar year y, we calculate the average default probability $\overline{\pi}^{P}_{y,aT}(d)$ and we then calculate the overall average default probability for rating a and maturity T, $\overline{\pi}^{P}_{aT}(d)$, by computing the mean across the N years, $\overline{\pi}^{P}_{aT}(d) = \frac{1}{N}\sum_{y=1}^{N}\overline{\pi}^{P}_{y,aT}(d)$. We denote by $\hat{\pi}^{P}_{aT}$ the corresponding historical default frequency given by Moody' for the period 1920–2012. For all major ratings (AAA, AA, A, BBB, BB, B, and C) and horizons of 1–20 years (Moody's only reports default rates for up to a horizon of 20 years) we find the value of d that minimizes the sum of absolute differences between the annualized historical and model-implied default rates by solving

$$\min_{\{d\}} \sum_{a=AAA}^{C} \sum_{T=1}^{20} \frac{1}{T} \left| \overline{\pi}^{P}_{aT}(d) - \hat{\pi}^{P}_{aT} \right|. \tag{7}$$

Using this approach our estimate is $\hat{d} = 0.8944$.

3.3 The term structure of default probabilities with the estimated default boundary

Our estimate of d matches the average default probability of firms issuing straight coupon bullet bonds to average historical default rates and there may be at least three concerns with this approach. First, firms issuing straight coupon bullet bonds may be different from the average firm in Moody's sample. Second, although the average historical default rate across maturities is matched, the term structure of default rates might not be matched accurately. Third, there may be systematic differences in the ability of the model to match default rates across ratings.

To address these concerns, we use the estimated default boundary to compute the average default probabilities for the broader sample of *all* rated industrial firms in Compustat and compare these to Moody's historical default rates. Specifically, we extract from Capital IG the issuer senior debt rating assigned by Standard & Poor's. There are almost no rating observations before 1985, so our sample period is 1985–2012. Table 7 gives summary statistics on the main parameters for this new sample. Compared to the sample of firms used to estimate the default boundary (in Table 6) the sample is 4.6 times as large and has a reasonably large number of speculative-grade firms.

We compute average model default probabilities for a given rating a in a manner analogous to the way Moody's calculates historical default rates. Specifically, on the final day of each month we find all firms for which we have data and that have rating a, and for those firms and dates we calculate the term structure of default probabilities, that is, default probabilities for horizons of 1, 2, . . ., 20 years. We then calculate the average term structure of default probabilities for each year 1985, . . ., 2012, and, finally, we compute the average

Table 7
Firm summary statistics, industrial firms in Compustat with a rating

	#Firms	Mean	10th	25th	Median	75th	90th
			Leverage ratio				
AAA	19	0.12	0.02	0.03	0.07	0.15	0.34
AA	95	0.15	0.04	0.07	0.11	0.18	0.29
A	385	0.20	0.06	0.10	0.17	0.27	0.39
BBB	650	0.28	0.08	0.15	0.25	0.37	0.51
BB	998	0.39	0.13	0.23	0.37	0.54	0.70
B	1,014	0.53	0.21	0.35	0.53	0.72	0.86
C	162	0.72	0.38	0.58	0.77	0.89	0.95
All	2,087	0.35	0.08	0.16	0.30	0.50	0.71
			Equity volatility				
AAA	19	0.26	0.17	0.21	0.26	0.30	0.35
AA	95	0.28	0.19	0.22	0.27	0.32	0.37
A	385	0.31	0.21	0.25	0.30	0.36	0.42
BBB	650	0.37	0.24	0.28	0.34	0.42	0.52
BB	998	0.48	0.31	0.37	0.45	0.56	0.68
B	1,014	0.65	0.38	0.48	0.61	0.77	0.94
C	162	0.88	0.53	0.62	0.80	1.05	1.24
All	2,087	0.45	0.24	0.30	0.40	0.55	0.72
			Asset volatility				
AAA	19	0.23	0.18	0.22	0.23	0.26	0.26
AA	95	0.24	0.20	0.21	0.24	0.26	0.29
A	385	0.26	0.20	0.22	0.24	0.28	0.35
BBB	650	0.28	0.20	0.22	0.27	0.32	0.38
BB	998	0.32	0.21	0.25	0.30	0.38	0.44
B	1,014	0.34	0.20	0.25	0.32	0.41	0.51
C	162	0.33	0.18	0.25	0.31	0.40	0.50
All	2,087	0.30	0.20	0.23	0.28	0.35	0.43
			Payout ratio				
AAA	19	0.027	0.008	0.014	0.023	0.035	0.050
AA	95	0.026	0.005	0.011	0.020	0.034	0.053
A	385	0.032	0.007	0.014	0.026	0.043	0.067
BBB	650	0.037	0.009	0.016	0.030	0.049	0.078
BB	998	0.039	0.009	0.018	0.033	0.054	0.078
B	1,014	0.049	0.010	0.024	0.044	0.068	0.092
C	162	0.067	0.022	0.045	0.068	0.089	0.108
All	2,087	0.039	0.008	0.017	0.032	0.054	0.081

For each firm in Compustat for which there is an S&P rating in Capital IQ, the leverage ratio, equity volatility, asset volatility, and payout ratio are calculated on December 31 in each year from 1985 to 2012. Leverage ratio is the ratio of the book value of debt to the market value of equity plus the book value of debt. Equity volatility is the annualized volatility of daily equity returns from the last 3 years. Asset volatility is the unlevered equity volatility, calculated as explained in the text. Payout ratio is yearly interest payments plus dividends plus share repurchases divided by firm value. Firm variables are computed using data from CRSP and Compustat.

term structure across years. We use the 10-year Treasury CMT rate as the risk-free rate because we do not have swap rates in the first years.

Figure 3 and Table 8 show the average model-implied default probabilities and Moody's historical default rates for 1920–2012. In both the figure and table we show 95% confidence bands for the historical default rate. The confidence band is obtained by following the simulation procedure in Section 2, where we simulate over 92 years and use Moody's historical default rates for the period 1920–2012 as input. In Table 8, cases where the model-implied default

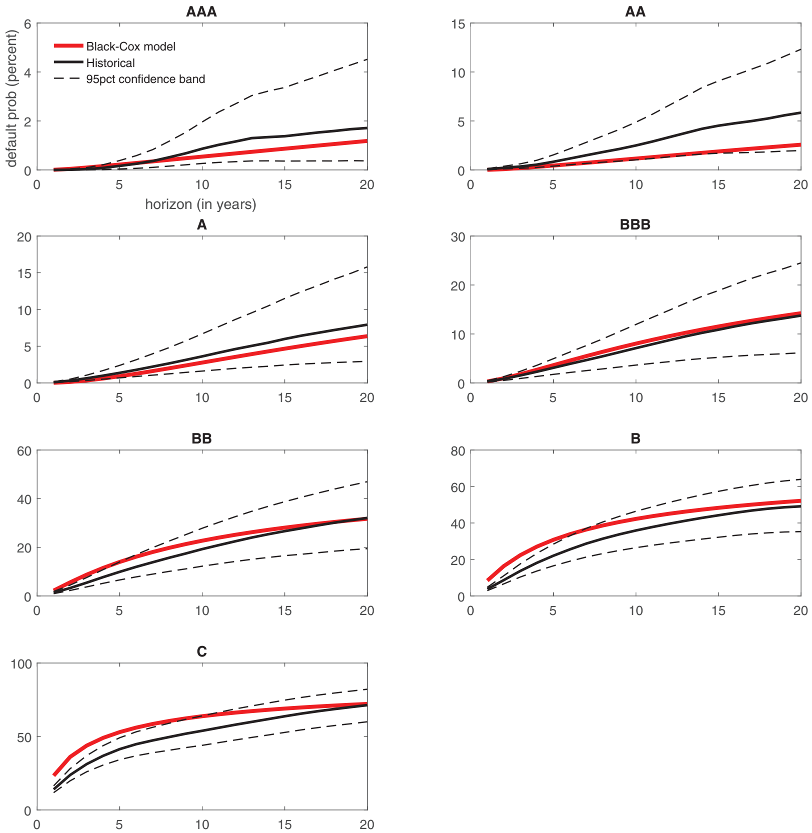

Figure 3
Average default probabilities in the Black-Cox model and historical default rates
We merge firm data from CRSP/Compustat with ratings from Standard & Poors, and, for every firm and every year from 1985 to 2012, we calculate a 1-, 2-, . . ., 19-, 20-year default probability in the Black-Cox model. The figure shows the average default probabilities along with the average historical default rate 1920–2012 calculated by Moody's. A 95% confidence band for the historical default rates are calculated following the approach in Section 2.1.

probability is outside the 95% and 99% confidence bands are indicated by * and **, respectively.

We see that in the Black-Cox model there is a statistically significant underestimation of 1- to 2-year AA and A default probabilities and overestimation of short-term speculative-grade default probabilities. The term structure of model-implied default probabilities is close to historical default rates for A beyond 3 years and for BBB-rated firms. These rating categories account for more than half of the U.S. corporate bond market volume (measured by the number of transactions).[12]

[12] According to TRACE Fact Book 2012, 53% of all U.S. corporate bond transaction volume in 2012 was in A- or BBB-rated bonds (tables C24 and C25).

Table 8
Average default probabilities in the Black-Cox model and historical default rates

Horizon (years)	1	2	3	4	5	6	8	10	12	15	20
AAA											
Model	0.01	0.04*	0.10*	0.16	0.22	0.28	0.41	0.54	0.67	0.87	1.18
Actual	0.00	0.01	0.03	0.09	0.17	0.25	0.52	0.87	1.16	1.38	1.71
95% c.b.	(NaN;NaN)	(0.00;0.04)	(0.00;0.09)	(0.01;0.21)	(0.04;0.38)	(0.06;0.58)	(0.16;1.17)	(0.26;1.96)	(0.34;2.69)	(0.36;3.37)	(0.37;4.52)
AA											
Model	0.01**	0.08**	0.19	0.31	0.45	0.59	0.87	1.16	1.45	1.89	2.58
Actual	0.07	0.22	0.35	0.54	0.83	1.17	1.83	2.50	3.34	4.52	5.85
95% c.b.	(0.03;0.13)	(0.11;0.37)	(0.17;0.62)	(0.25;0.99)	(0.39;1.54)	(0.53;2.17)	(0.79;3.48)	(1.03;4.88)	(1.34;6.59)	(1.70;9.09)	(2.0;12.3)
A											
Model	0.02**	0.14**	0.34	0.60	0.91	1.26	2.00	2.77	3.55	4.67	6.37
Actual	0.10	0.31	0.64	0.99	1.38	1.78	2.66	3.62	4.61	5.99	7.93
95% c.b.	(0.05;0.17)	(0.17;0.51)	(0.34;1.06)	(0.51;1.69)	(0.70;2.39)	(0.88;3.16)	(1.24;4.82)	(1.62;6.69)	(1.99;8.64)	(2.4;11.5)	(2.9;15.8)
BBB											
Model	0.30	0.95	1.78	2.70	3.64	4.57	6.37	8.04	9.54	11.53	14.24
Actual	0.29	0.86	1.54	2.29	3.10	3.90	5.46	7.11	8.72	10.87	13.76
95% c.b.	(0.18;0.44)	(0.52;1.29)	(0.92;2.38)	(1.33;3.60)	(1.75;4.95)	(2.14;6.31)	(2.87;9.01)	(3.6;11.9)	(4.3;14.8)	(5.2;18.9)	(6.2;24.5)
BB											
Model	2.23**	5.61**	8.76**	11.54*	13.99	16.15	19.78	22.72	25.15	28.09	31.74
Actual	1.35	3.27	5.46	7.75	9.92	11.97	15.71	19.27	22.47	26.65	32.06
95% c.b.	(0.95;1.83)	(2.27;4.48)	(3.73;7.58)	(5.2;10.8)	(6.6;14.0)	(7.8;17.0)	(10.1;22.6)	(12.2;27.8)	(14.2;32.6)	(16.6;38.8)	(19.4;47.0)
B											
Model	8.43**	16.40**	22.39**	27.04**	30.78**	33.86*	38.65	42.24	45.05	48.31	52.17
Actual	3.80	8.71	13.72	18.16	22.06	25.54	31.41	35.89	39.58	44.22	49.14
95% c.b.	(2.93;4.78)	(6.6;11.1)	(10.4;17.6)	(13.7;23.3)	(16.5;28.4)	(19.0;33.0)	(23.3;40.6)	(26.5;46.4)	(29.0;51.3)	(32.2;57.3)	(35.3;63.9)
C											
Model	23.32**	36.11**	43.79**	49.07**	52.99**	56.05**	60.57*	63.80	66.24	69.00	72.17
Actual	14.02	23.81	31.21	36.86	41.40	44.78	49.63	53.88	58.02	63.76	71.34
95% c.b.	(11.8;16.4)	(19.9;28.1)	(26.0;37.0)	(30.6;43.8)	(34.2;49.1)	(36.9;53.2)	(40.6;59.0)	(44.0;64.0)	(47.6;68.8)	(52.8;74.8)	(60.0;82.2)

We merge firm data from CRSP/Compustat with ratings from Standard & Poors, and, for every firm and every year from 1985 to 2012, we calculate a 1-, 2-, ..., 19-, 20-year default probability in the Black-Cox model. ‘Model’ shows the average default probabilities. ‘Actual’ shows Moody’s average historical default rates from 1920 to 2012. ‘95% c.b.’ shows 95% confidence bands for the historical default rates calculated following the approach in Section 2.1 . * and ** show when the model-implied default probability is outside the 95% and 99% confidence band, respectively.

It may seem surprising that the Black-Cox model captures the term structure of default rates for BBB-rated bonds well because there are a number of papers showing that, for horizons less than 3–4 years, structural models imply essentially zero default probabilities for investment-grade firms (Zhou 2001; Leland 2004, 2006; Cremers et al. 2008; Zhang et al. 2009; and others). We show in Appendix C that results in the existing literature documenting a failure of structural models to capture short-term default rates are strongly biased due to a "convexity effect" arising from Jensen's inequality. The bias arises when using a representative firm with average leverage (and average asset volatility and payout rate) to calculate short-term default probabilities because the default probability using average leverage is substantially lower than the average default probability calculated using the distribution of leverage. Our results show that once we deal with the convexity bias by using data on individual firms the Black-Cox model captures short-term default rates much better than previously reported.

3.4 Average corporate bond credit spreads

We calculate average spreads by following the calculations in Duffee (1998). Specifically, we calculate a monthly average actual spread for a given rating a, maturity range M_1 to M_2, and month t. To ease notation, we index the combination of rating, maturity range and month by $h=(a,[M_1;M_2],t)$. For a given h, we find all N_h bond observations with rating a and individual bond maturities $T_1^h,T_2^h,\ldots,T_{N_h}^h$, where $M_1 \leq T_i^h < M_2$, observed on days $\tau_1^h,\tau_2^h,\ldots,\tau_{N_h}^h$ in month t. Denoting the corresponding yield observations as $y_1^h,y_2^h,\ldots,y_{N_h}^h$ and the swap rates as $sw(\tau_1^h,T_1^h),sw(\tau_2^h,T_2^h),\ldots,sw(\tau_{N_h}^h,T_{N_h}^h)$, the average yield spread for rating a, maturity range M_1 to M_2, and month t is

$$s^h=\frac{1}{N^h}\sum_{i=1}^{N^h}\left(y_i^h-sw(\tau_i^h,T_i^h)\right). \tag{8}$$

The average yield spread for a given rating and a given maturity interval is then the average of the monthly values.

Similarly, we calculate model-implied spreads by replacing the actual spread $y_i^h-sw(\tau_i^h,T_i^h)$ with the spread $s(\hat{d},\Theta^Q_{j_i^h},L_{j_i^h,\tau_i^h},T_i^h)$ implied by the Black-Cox model given in Equation (1), where $\Theta^Q_{j_i^h}=(sw(\tau_i^h,T_i^h),\sigma_{j_i^h},\delta_{j_i^h,\tau_i^h},R)$, $\sigma_{j_i^h}$ is the asset volatility of firm j that issued bond i, $L_{j_i^h,\tau_i^h}$ and $\delta_{j_i^h,\tau_i^h}$ are the leverage ratio and payout rate, respectively, of firm j on day τ_i^h, $R=37.8\%$ is the recovery rate, and the default boundary $\hat{d}=0.8944$, as estimated in Section 3.2.

We compute confidence intervals for the model-implied spreads in the following way. In Section 2.2 we calculate the distribution of the default boundary. We redo this calculation using Moody's default rates from 1920 to 2012. That is, we simulate over 92 years and set the leverage ratios for each rating such that the historical 10-year cumulative default rates for

AAA, AA, ..., C for the period 1920–2012 are matched and use for each rating the average cohort size for the period 1970–2012 in our simulation. For the estimated default boundary distribution, we calculate the 2.5% and 97.5% quantiles, $q^{2.5\%}$ and $q^{97.5\%}$ and, since the default boundary is scale independent and is set to unity in the simulations, we compute the 2.5% and 97.5% quantiles for $\hat{d}$ as $q^{2.5\%}\hat{d}$ and $q^{97.5\%}\hat{d}$, respectively. Finally, since the model-implied spread is monotone in the default boundary, we calculate the corresponding quantiles for the model-implied spread using these same values ($q^{2.5\%}\hat{d}$ and $q^{97.5\%}\hat{d}$).

3.4.1 Sorting by rating. Table 9 shows actual and model-implied bond spreads in our sample, calculated using the approach just described. We see that the average actual investment-grade spread across maturity is 92 bps, whereas the average model-implied spread is 111 bps. Accounting for the uncertainty of default probabilities, the difference is statistically insignificant. We also see good correspondence between model-implied and actual investment-grade spreads when we look at the individual maturities 3–7, 7–13, and 13–20 years. Thus, the Black-Cox model captures average investment-grade spreads well. Turning to speculative-grade spreads, we see an underprediction of spreads across maturity with average actual spreads at 544 bps and average model spreads at 382 bps with the difference being statistically significant.

When we look at individual investment-grade ratings in Table 9 the Black-Cox model slightly underpredicts spreads on AAA- and AA-rated bonds (by 13–15 bps) and overpredicts spreads on A-rated bonds by 24 bps. For A-rated bonds the average actual spread across maturity is 61 bps, whereas it is 85 bps in the model. The average actual BBB spread across maturity is 146 bps, whereas the average model-implied spread is 169 bps. Thus, the average spread of bonds where most trading takes place in the U.S corporate bonds—bonds with a rating of A or BBB—is captured well by the Black-Cox model. For speculative-grade bonds, underprediction increases as we move down the rating scale. For BB-rated bonds, there is no significant underprediction for maturities below 13 years, while model spreads are too low for longer maturities.

Overall, the average level of actual and model-implied investment-grade spreads is statistically not different. In contrast we find that the model underpredicts speculative-grade spreads; here, the average spread across maturity is 544 bps in the data and 382 bps in the model, and the difference is statistically significant.

3.4.2 Sorting by yield spread. The literature has traditionally compared model-implied and actual credit spreads within rating categories. There are several reasons for this. First, Moody's provides yield data and default rates from 1920, and there is therefore a long history of default and yield organized by rating. In fact, as far as we are aware, the only publicly available data on aggregate default rates are organized by rating. Second, spreads organized by rating show a large variation in the mean, with lower-rated firms having

Table 9
Actual and model yield spreads

		3–20y	3–7y	7–13y	13–20y
Inv	Actual spread	92	89	87	87
	Model spread	111	107	107	88
		(88;128)	(82;127)	(87;121)	(76;96)
	Difference	19	18	20*	1
		(−4;36)	(−8;38)	(0;34)	(−11;9)
	Observations	294	294	293	244
Spec	Actual spread	544	560	417	461
	Model spread	382	376	392	314
		(305;440)	(289;443)	(336;429)	(279;337)
	Difference	−162**	−184**	−25	−147**
		(−239;−104)	(−271;−117)	(−81;12)	(−182;−124)
	Observations	289	276	229	141
AAA	Actual spread	16	4	6	22
	Model spread	3	3	1	2
		(2;4)	(1;6)	(0;1)	(1;2)
	Difference	−13**	−0	−6**	−20**
		(−14;−12)	(−3;2)	(−6;−5)	(−20;−20)
	Observations	132	70	70	91
AA	Actual spread	23	17	34	26
	Model spread	9	2	14	19
		(6;10)	(1;3)	(11;17)	(15;22)
	Difference	−15**	−15**	−20**	−7**
		(−17;−13)	(−16;−14)	(−23;−17)	(−11;−4)
	Observations	289	279	264	93
A	Actual spread	61	50	65	63
	Model spread	85	67	102	83
		(67;99)	(48;82)	(82;117)	(71;90)
	Difference	24**	17	37**	19**
		(6;38)	(−2;32)	(17;51)	(8;27)
	Observations	294	294	293	223
BBB	Actual spread	146	141	141	144
	Model spread	169	165	166	131
		(134;195)	(126;195)	(137;186)	(112;143)
	Difference	23	24	25	−14*
		(−12;49)	(−15;54)	(−4;46)	(−32;−1)
	Observations	291	291	257	198
BB	Actual spread	377	370	290	398
	Model spread	349	320	337	255
		(282;397)	(247;374)	(285;372)	(223;277)
	Difference	−27	−51	46	−142**
		(−94;21)	(−124;4)	(−5;82)	(−175;−121)
	Observations	259	240	216	114
B	Actual spread	675	723	427	445
	Model spread	445	480	441	323
		(360;509)	(376;560)	(376;485)	(294;342)
	Difference	−229**	−243**	15	−122**
		(−314;−166)	(−347;−163)	(−51;59)	(−150;−103)
	Observations	243	203	134	82
C	Actual spread	1,442	1,211	1,948	661
	Model spread	958	1,097	783	525
		(828;1,041)	(920;1,209)	(709;829)	(449;575)
	Difference	−484**	−114*	−1,165**	−136**
		(−615;−401)	(−291;−2)	(−1239;−1,119)	(−212;−86)
	Observations	96	65	42	7

This table shows actual and model-implied corporate bond yield spreads. Spreads are grouped by remaining bond maturity at the quotation date. 'Actual spread' is the average actual spread to the swap rate. "Model spread" is the average Black-Cox model spreads of the bonds in a given maturity/rating bucket. The average spread is calculated by first calculating the average spread of bonds in a given month and then calculating the average of these spreads over months. "Difference" is the difference between the model spread and the actual spread. In parentheses are 95% confidence bands calculated according to Section 2.2; * implies significance at the 5% level and ** at the 1% level. "Observations" is the number of monthly observations. The bond yield spreads are from the period 1987–2012.

higher average spreads; matching average bond spreads organized by rating has provided a hard test for structural models.

Although average default rates are available only by rating, we can nevertheless sort bonds in other ways in order to compare model-implied and actual spreads. If there is a substantial difference, the model is misspecified in some dimension. Since any useful sort should result in significant variation in spreads, the most obvious choice is to sort according by actual spreads.

Table 10 shows model spreads sorted by the size of the actual spread. We see that for actual spreads below 1,000 bps there is no statistical difference between model spreads and actual spreads when averaged across maturity. For example, the actual spread for bonds with spreads between 100 and 150 bps is 121 bps, whereas it is 136 bps in the model. However, for bond spreads above 150 bps, we start to see a modest underprediction at long maturities, and it becomes strong only when spreads are above 300 bps. However, above 1,000 bps the model substantially underestimates spreads, and here the average model spread is only around half of the average actual spread.

Overall, the results when sorting by actual spread are similar to those sorted by rating, namely that spreads for low credit risk firms are matched well while spreads for the highest credit risk firms, particular for bonds with long maturity, are under-predicted.

3.5 Time-series variation in yield spreads

Having established that the Black-Cox model can match the average size of investment-grade credit spreads, we next examine whether the model can also capture their time-series variation. In each month, we calculate the average actual yield spread for a given rating according to Equation (8) (where the spread is relative to the swap rate) along with the corresponding model-implied average spread and investigate the monthly time series.

To provide an overall assessment of the model's ability to capture investment-grade spreads, we group together all investment-grade spreads and all maturities between 3 and 20 years and plot the actual and model-implied spreads in Figure 4. We see that the model-implied spread tracks the actual spread well with a correlation is 93%.

To test more formally the ability of the Black-Cox model to capture the time-series variation in spreads, we regress the monthly time series of the actual spread, s_t, on the model-implied spread, $\hat{s}_t$,

$$s_t = \alpha + \beta \hat{s}_t + \epsilon_t, \tag{9}$$

and report the β and the R^2 of the regression in Table 11, panel A. The table shows that for all bonds with maturities between 3 and 20 years the regression of actual investment-grade spreads on model-implied investment-grade spreads gives a slope coefficient of 0.88 and an R^2 equal to 87% showing that once investment-grade spreads are aggregated model-implied spreads track actual spreads very well. The R^2's for the separate aggregate regressions for A and

Table 10
Actual and model yield spreads sorted by actual spread

		3–20y	3–7y	7–13y	13–20y
<20bps	Actual spread	7	7	8	10
	Model spread	15	14	18	19
		(11;18)	(9;18)	(14;22)	(16;22)
	Difference	8**	7**	10**	9**
		(3;11)	(2;11)	(6;13)	(5;12)
	Observations	279	279	214	138
20–40bps	Actual spread	30	29	30	31
	Model spread	37	30	54	36
		(27;45)	(20;39)	(41;63)	(30;41)
	Difference	7	1	23**	5
		(−3;15)	(−9;10)	(11;32)	(−1;9)
	Observations	279	272	233	165
40–70bps	Actual spread	55	54	55	56
	Model spread	72	69	84	56
		(55;85)	(49;85)	(66;98)	(47;62)
	Difference	17	15	29**	0
		(−0;30)	(−5;31)	(11;42)	(−9;7)
	Observations	284	277	262	191
70–100bps	Actual spread	84	85	84	84
	Model spread	104	115	118	87
		(80;122)	(81;142)	(94;135)	(74;96)
	Difference	20	30	34**	3
		(−4;38)	(−4;57)	(10;50)	(−10;12)
	Observations	281	264	246	170
100–150bps	Actual spread	121	121	120	121
	Model spread	136	145	140	121
		(105;160)	(102;179)	(112;160)	(104;133)
	Difference	15	25	19	−0
		(−16;39)	(−19;59)	(−9;39)	(−18;11)
	Observations	269	260	254	166
150–200bps	Actual spread	172	172	172	171
	Model spread	174	167	188	144
		(135;203)	(120;204)	(153;213)	(123;158)
	Difference	2	−5	15	−27**
		(−37;31)	(−52;32)	(−20;40)	(−48;−13)
	Observations	252	222	211	120
200–300bps	Actual spread	243	242	243	245
	Model spread	257	251	305	218
		(202;297)	(185;303)	(253;343)	(188;238)
	Difference	14	9	62*	−28**
		(−40;54)	(−57;60)	(10;100)	(−57;−8)
	Observations	267	222	220	99
300–1,000bps	Actual spread	499	523	456	517
	Model spread	507	558	499	368
		(414;573)	(442;642)	(430;546)	(332;391)
	Difference	8	35	43	−149**
		(−85;74)	(−81;119)	(−26;90)	(−185;−127)
	Observations	268	244	221	150
>1,000bps	Actual spread	1,744	1,746	1,789	1,166
	Model spread	909	1,004	735	513
		(747;1,026)	(806;1,151)	(646;792)	(455;550)
	Difference	−835**	−742**	−1,055**	−653**
		(−997;−718)	(−940;−595)	(−1,143;−998)	(−711;−616)
	Observations	132	89	59	15

This table shows actual and model-implied corporate bond yield spreads. Spreads are grouped by the size of the actual spread and the remaining bond maturity at the quotation date. 'Actual spread' is the average actual spread to the swap rate. 'Model spread' is the average Black-Cox model spreads of the bonds in a given maturity/rating bucket. The average spread is calculated by first calculating the average spread of bonds in a given month and then calculating the average of these spreads over months. 'Difference' is the difference between the model spread and the actual spread. In parentheses are 95% confidence bands calculated according to Section 2.2; * implies significance at the 5% level and ** at the 1% level. 'Observations' is the number of monthly observations. The bond yield spreads are from the period 1987–2012.

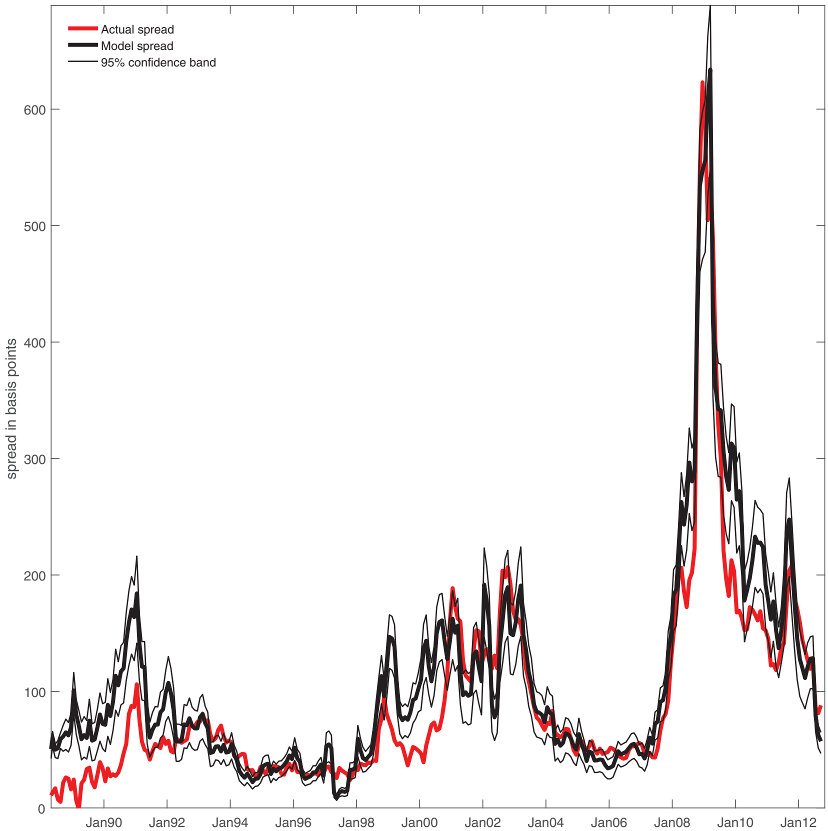

Figure 4
Time-series variation in investment-grade spreads
This graph shows the time series of actual and model-implied investment-grade corporate bond spreads. Each month all daily yield observations in bonds with an investment-grade rating and with a maturity between 3 and 30 years are collected and the average actual spread (to the swap rate) and the average model-implied spread in the Black-Cox model are computed. The graph shows the time series of monthly spreads. A 95% confidence band for the model-implied spread is calculated following the approach in Section 2.2.

BBB spreads are high at 70% and 88%, respectively. For speculative-grade and AAA/AA ratings, the ability of the Black-Cox model to capture the time-series variation is much lower. More noise due to fewer observations is one factor contributing to the deteriorating fit.

Panel B shows the regression in changes,

$$s_{t+1}-s_t=\alpha+\beta(\hat{s}_{t+1}-\hat{s}_t)+\epsilon_{t+1}, \tag{10}$$

and we see that the R^2's are substantially lower. This implies that significant variation in monthly changes in credit spreads is not explained by the Black-Cox

Table 11
Commonality in time-series variation of actual and model-implied yield spreads

		3–20y	3–7y	7–13y	13–20y
A. Regression in levels					
Inv	β	0.88 (0.06)	0.81* (0.08)	0.82 (0.09)	0.79 (0.13)
	R^2	0.87	0.84	0.69	0.59
Spec	β	0.85 (0.38)	0.82 (0.41)	0.80 (0.19)	1.55** (0.20)
	R^2	0.16	0.15	0.25	0.75
AAA	β	1.17 (0.44)		1.40** (0.14)	1.63** (0.22)
	R^2	0.15		0.26	0.22
AA	β	1.07 (0.58)	0.81 (0.34)	1.57* (0.22)	−0.10** (0.04)
	R^2	0.22	0.03	0.62	0.08
A	β	0.65** (0.08)	0.52* (0.19)	0.54** (0.11)	0.40** (0.09)
	R^2	0.70	0.47	0.57	0.41
BBB	β	0.82** (0.07)	0.71** (0.09)	0.90 (0.10)	0.72 (0.15)
	R^2	0.88	0.81	0.74	0.62
BB	β	0.72 (0.28)	0.70 (0.31)	0.64* (0.14)	1.66* (0.28)
	R^2	0.42	0.39	0.49	0.79
B	β	0.61 (0.28)	0.64 (0.30)	0.31** (0.19)	1.14 (0.23)
	R^2	0.09	0.10	0.06	0.58
C	β	−0.75** (0.56)	−0.34* (0.56)		
	R^2	0.04	0.01		
B. Regression in changes					
Inv	β	0.52** (0.08)	0.47** (0.08)	0.37** (0.09)	0.62** (0.10)
	R^2	0.35	0.33	0.17	0.37
Spec	β	0.00* (0.50)	0.42 (0.49)	0.76 (0.22)	1.30 (0.24)
	R^2	0.00	0.01	0.17	0.45
AAA	β	0.56 (0.27)		0.54 (0.49)	0.57 (0.41)
	R^2	0.11		0.07	0.08
AA	β	0.25** (0.11)	0.21** (0.30)	0.23** (0.12)	−0.08** (0.05)
	R^2	0.06	0.01	0.06	0.11
A	β	0.37** (0.07)	0.50** (0.07)	0.18** (0.06)	0.47** (0.08)
	R^2	0.27	0.44	0.10	0.36
BBB	β	0.46** (0.08)	0.33** (0.08)	0.31** (0.11)	0.67** (0.09)
	R^2	0.32	0.20	0.11	0.54
BB	β	0.66** (0.11)	0.61** (0.12)	0.68** (0.10)	1.33 (0.30)
	R^2	0.34	0.31	0.47	0.41
B	β	−1.02** (0.32)	−0.92** (0.33)	−0.18** (0.26)	1.39 (0.37)
	R^2	0.14	0.13	0.01	0.41
C	β	−2.12* (1.49)	−2.37 (2.28)		
	R^2	0.08	0.06		

For a given rating and maturity group we calculate a monthly average spread by computing the average yield spread for bonds with the corresponding rating and maturity observed in that month. We do this for both model-implied spreads and actual spreads (to the swap rate) resulting in a time series of monthly actual spreads $s_1, s_2, \ldots, s_T$ and implied spreads from the Black-Cox model $\hat{s}_1, \hat{s}_2, \ldots, \hat{s}_T$ for the period 1987-2012. Panel A shows the regression coefficient in the regression of the actual spread on the model-implied spread $s_t = \alpha + \beta \hat{s}_t + \epsilon_t$. In parentheses is the standard error, Newey-West corrected with 12 lags and * implies that β is significantly different from one at the 5% level and ** at the 1% level. In some months there may not be any observations and if there are less than 100 monthly observations we do not report regression coefficients. Panel B shows regression results for monthly changes, $s_{t+1} - s_t = \alpha + \beta(\hat{s}_{t+1} - \hat{s}_t) + \epsilon_{t+1}$. In parentheses is the ordinary least squares (OLS) standard error, and * implies that β is significantly different from one at the 5% level and ** at the 1% level.

model and Collin-Dufresne, Goldstein, and Martin (2001) and Feldhütter (2012) link this variation to supply/demand shocks.[13]

3.6 Spread predictions on individual bonds

Our main result is that, when calibrated to match historical default rates, the Black-Cox model with a constant (1) Sharpe ratio, (2) recovery rate, and (3) default boundary, and no priced risks beyond diffusion risk can match the average spread of investment-grade bonds. This result does not necessarily imply that the model can match spreads on individual bonds with a high degree of precision, because average spreads may well mask significant individual pricing errors. While our interest lies mainly in asking whether the model can capture average spreads, we nevertheless carry out an exploratory analysis on the ability of the model to capture the cross-section of spreads (for a more extensive analysis, see Bao 2009).

The first column in Table 12 shows the R^2's from regressing actual spreads on model-implied spreads (and a constant) at the individual bond level. For investment-grade bonds the R^2 is 44%, which is substantially below the R^2 of 87% obtained using monthly average spreads and reported in Table 11. For speculative-grade bonds, the explanatory power of the regression at the individual bond level is only 13%, showing that the model has only limited ability to price speculative-grade bonds.

To give an indication on how the model or parameter estimates might be improved, we correlate the pricing error—the difference between the actual and model-implied spread—with variables used in the estimation. Table 12 shows the results. The pricing errors for investment-grade bonds have a correlation of -0.45 with leverage and -0.34 with the payout rate. This suggests that estimates for individual bonds could be improved either by estimating leverage and payout rate in a different way or indeed by changes to the model. Correlations between (equity and asset) volatilities and pricing errors are modest and range from -0.14 to 0.12. This may indicate that a better measurement of volatility and/or incorporation of stochastic asset volatility into the model may be less important in improving cross-sectional accuracy.

3.7 The role of bond illiquidity

A number of papers examine the impact of illiquidity on corporate bond spreads; these include Dick-Nielsen, Feldhütter, and Lando (2012) (DFL), Bao, Pan, and Wang (2011), Friewald, Jankowitsch, and Subrahmanyam (2012), and Lin, Wang, and Wu (2011). These papers examine transactions data from the relatively recent past, typically from 2004, but since we use spread data starting back in 1987, we can provide evidence on the impact of illiquidity on credit

[13] Results are similar when we sort by absolute spread change instead of rating (see the Internet Appendix for details).

Table 12
Explaining individual pricing errors

		Correlation of pricing error with			
	R^2	L_t	σ_t^e	σ_t^a	δ_t
Investment grade	0.44	−0.45	−0.14	0.12	−0.34
Speculative grade	0.13	−0.02	−0.03	−0.02	0.00

The first column shows the R^2 from running a regression of actual spreads on individual bonds on the implied spreads from the Black-Cox where we use all transactions in the data sample, separated into investment grade and speculative grade. The next columns show the correlation between the pricing error, defined as the difference between the actual spread and model-implied spread, and variables that may contribute to pricing errors. L_t is the leverage ratio on the day of the transaction, σ_t^e is the estimated equity volatility on the day of the transaction, σ_t^a is the issuing firm's asset volatility when estimated day-by-day, δ_t is the payout rate on the day of the transaction. The bond yield spreads are from the period 1987–2012.

Table 13
Credit spread residuals sorted on bond age

	(1)	(2)	(3)	(4)	(5)
Inv	**18**	**8**	**19**	**−5**	**30***
	(−4;35)	(−11;23)	(−5;37)	(−32;15)	(5;48)
	[28,320]	[26,855]	[26,546]	[30,998]	[31,125]
Spec	**88****	**56**	**−114****	**−356****	**−382****
	(34;125)	(−9;101)	(−203;−51)	(−424;−306)	(−430;−350)
	[7,425]	[8,940]	[9,321]	[4,865]	[4,782]

We first sort all bond spread observations for bonds with a maturity between 3 and 20 years into quintiles based on the time since the bond was issued. For investment- and speculative-grade bonds, we then calculate the average difference between the model-implied and actual spread. The table shows this average difference in basis points. In parentheses are 95% confidence bands calculated according to Section 2.2. * implies significance at the 5% level and ** at the 1% level. The number of observations are in brackets.

spreads over a longer historical time period. The drawback of our longer time period is that we cannot calculate transactions-based measures of illiquidity.

Instead we use bond age as a measure of bond illiquidity since age is known to be related to illiquidity (see Bao et al. 2011; Houweling et al. 2005; and the references therein). In Table 13 we first sort average credit spread residuals, defined as the difference between the model-implied and actual credit spread, into quintiles by bond age and then by whether the bond is investment grade or speculative grade. The table reports the average credit spread residual and tests whether this average is different from zero.

For investment-grade bonds there is essentially no relation between the average pricing error and bond illiquidity. This is consistent with the findings in DFL, Friewald, Jankowitsch, and Subrahmanyam (2012), and Lin, Wang, and Wu (2011) that the potential impact of illiquidity on prices of investment-grade bonds is much smaller than for speculative grade.

In contrast, Table 13 shows a strong relation between pricing errors and bond illiquidity for speculative-grade bonds. For liquid speculative-grade bonds the pricing error is modestly positive, but as we move to more illiquid bonds a strong model underprediction emerges and the average pricing error difference between bonds in the least and most liquid quintiles is 470 bps. The magnitude of the pricing error difference across bond illiquidity suggests that much of the underprediction of speculative-grade credit spreads can be explained by bond illiquidity. This in turn suggests incorporating illiquidity into structural models is important to price speculative-grade bonds. As an interesting class of such models, He and Milbradt (2014) and Chen et al. (2017) incorporate search frictions in a structural model of credit risk in such a way that illiquidity is more important for speculative-grade bonds.

For the 2008–2009 financial crisis, DFL find an illiquidity premium in AAA-, AA-, A-, and BBB-rated and speculative-grade bonds of 5, 42, 51, 93, and 197 bps, respectively. If we restrict the analysis in Section 3.4 to the crisis period identified in DFL (2007:Q2–2009:Q2), for maturities between 3 and 20 years we find that the average difference between actual and model-implied spreads for AAA-, AA-, A-, and BBB-rated and speculative-grade bonds is 18, 63, −43, −33, and 262 bps, respectively (when calculated like in Table 9). For AAA- and AA-rated bonds, the model underprediction is comparable to the liquidity premium found in DFL and for speculative-grade bonds the underprediction is around 100 bps larger. One has to be careful in interpreting point estimates of average spread differences in basis points over a relatively short period where spreads were at a historical high and very volatile. But taking the model overprediction of 43 and 33 bps for A- and BBB-rated bonds at face value suggests the presence of some model misspecification during the financial crisis.

3.8 Using default data from 1970 to 2012 to estimate the default boundary

We saw in Section 1 that when the historical BBB and AAA default rates are used one at a time as estimates of the BBB and AAA default probabilities, the appearance of a credit spread puzzle strongly depends on the time period over which historical default rates are calculated.

To see whether our proposed method suffers from the same drawback, we estimate the default boundary as described in Section 3.2 but fitting to Moody's default rates from 1970 to 2012 rather than from 1920 to 2012 like in the main analysis. In this case the default boundary is estimated to be $\hat{d}=0.9302$ (compared with 0.8944 found for 1920–2012) and Table 14 shows average spreads using this value. The table shows that there are only modest changes in the model-implied spreads. For example, the average investment- and speculative-grade spreads are 122 and 420 bps, respectively, when using default rates from 1970 to 2012 compared to 111 and 382 bps, respectively, when using default rates from 1920 to 2012.

Figure 5 shows the results, using our proposed approach, in the same format as Figure 1 calibrated both to 1970–2012 and to 1920–2012. Unlike the earlier results we see that results are very similar. Indeed, given the differences we observe in Figure 1, the stability of the model-implied spreads is striking and suggests that by using a cross-section of default rates to calibrate the model, we can provide both firmer and more stable conclusions.

Table 14
Actual and model yield spreads when using default rates from 1970 to 2012 to calibrate the model

		3–20y	3–7y	7–13y	13–20y
Inv	Actual spread	92	89	87	87
	Model spread	122 (91;150)	120 (85;154)	116 (90;140)	93 (78;106)
	Difference	30 (−1;59)	31 (−4;65)	30* (3;53)	7 (−9;20)
	Observations	294	294	293	244
Spec	Actual spread	544	560	417	461
	Model spread	420 (315;517)	420 (301;535)	417 (344;476)	329 (284;364)
	Difference	−124** (−229;−27)	−140** (−259;−24)	−0 (−73;59)	−132** (−177;−97)
	Observations	289	276	229	141
AAA	Actual spread	16	4	6	22
	Model spread	4 (2;6)	5 (1;11)	1 (0;1)	2 (2;3)
	Difference	−12** (−14;−10)	1 (−2;7)	−6** (−6;−5)	−20** (−20;−19)
	Observations	132	70	70	91
AA	Actual spread	23	17	34	26
	Model spread	10 (7;13)	2 (1;4)	16 (11;21)	21 (16;26)
	Difference	−14** (−17;−11)	−14** (−15;−13)	−18** (−23;−13)	−5 (−10;0)
	Observations	289	279	264	93
A	Actual spread	61	50	65	63
	Model spread	94 (69;117)	77 (51;103)	112 (85;136)	88 (73;100)
	Difference	33** (8;56)	27* (1;53)	46** (19;70)	24** (10;37)
	Observations	294	294	293	223
BBB	Actual spread	146	141	141	144
	Model spread	186 (139;230)	185 (131;236)	180 (141;213)	139 (115;158)
	Difference	40 (−7;83)	43 (−10;94)	39* (0;72)	−6 (−29;14)
	Observations	291	291	257	198
BB	Actual spread	377	370	290	398
	Model spread	381 (291;460)	356 (256;448)	360 (292;416)	270 (228;303)
	Difference	5 (−85;84)	−15 (−114;78)	70* (2;126)	−128** (−170;−95)
	Observations	259	240	216	114
B	Actual spread	675	723	427	445
	Model spread	487 (371;594)	532 (390;670)	471 (385;540)	336 (299;364)
	Difference	−187** (−303;−80)	−190** (−333;−53)	44 (−42;113)	−109** (−146;−81)
	Observations	243	203	134	82
C	Actual spread	1,442	1,211	1,948	661
	Model spread	1,014 (846;1,136)	1,173 (946;1,339)	814 (720;881)	559 (460;635)
	Difference	−429** (−596;−306)	−39 (−265;128)	−1,134** (−1,228;−1,066)	−102** (−201;−25)
	Observations	96	65	42	7

In the main analysis the default boundary is estimated using Moody's default rates from 1920 to 2012. This table shows results when the default boundary is estimated using Moody's default rates from 1970 to 2012. The table shows actual and model-implied corporate bond yield spreads. Spreads are grouped by remaining bond maturity at the quotation date. 'Actual spread' is the average actual spread to the swap rate. 'Model spread' is the average Black-Cox model spreads of the bonds in a given maturity/rating bucket. The average spread is calculated by first calculating the average spread of bonds in a given month and then calculating the average of these spreads over months. 'Difference' is the difference between the model spread and the actual spread. In parentheses are 95% confidence bands calculated according to Section 2.2. * implies significance at the 5% level and ** at the 1% level. 'Observations' is the number of monthly observations. The bond yield spreads are from the period 1987–2012.

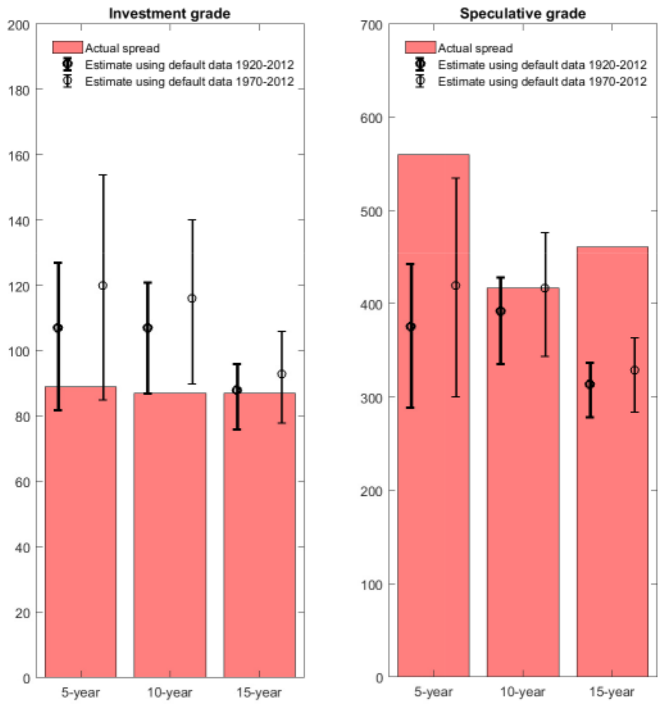

Figure 5
Actual and model-implied corporate bond yield spreads when using default rates from 1970 to 2012 and 1920 to 2012
This figure shows average actual and model-implied corporate bond yield spreads estimated using default rates from either 1970 to 2012 or 1920 to 2012. Model-implied spreads are calculated according to our proposed method where many default rates across maturity and rating are used in the calibration of the model and the figure shows results when default rates from either 1970–2012 or 1920–2012 are used in the calibration. Confidence bands take into account uncertainty about ex ante default probabilities. Spreads are from Tables 9 and 14. Actual bond yield spreads are average spreads to the swap rate from noncallable bonds issued by industrial firms and from the period 1987–2012.

3.9 The relation between the estimated default boundary and existing estimates in the literature

In Section 3.2 we estimate the default boundary to be $\hat{d} = 0.8944$; that is, a firm defaults when its (firm) value is less than 89.44% of the face value of debt. Davydenko (2013) studies the location of the default boundary, measured as the total market value of the firm in the month preceding default expressed as a fraction of the face value of debt. He finds that the average default boundary is 66.0%.

Although our estimate appears to be substantially higher, note that Davydenko's boundary is measured as firm value relative to the face value of debt *at the time of default*. Ours, on the other hand, is measured relative to the face value of debt *at the time we observe a bond price*. To see whether this difference in definition may explain the gap between our estimate of d

and Davydenko's, we use our estimate to calculate a rough estimate of d under Davydenko's definition.

The average bond maturity in our sample is 5.72 years, that is, close to 6 years. We use Moody's default database and find all defaults over the period 1990–2012 for which we can identify the following three data items: (1) the face value of debt in the year prior to default; (2) the face value of debt 7 years prior to default (i.e., 6 years earlier); and (3) a rating 7 years prior to default. We have 128 such observations.

We find the average log growth rate in the face value of debt from 7 years prior to default to 1 year prior to default to be 27.76%. Thus, based on our estimate of the default boundary of 89.44% and the growth rate in the face value of debt, the average default boundary in terms of Davydenko's definition is $0.8944e^{-0.2776} = 67.76\%$, close to Davydenko's estimates of 66.0%.

The fact that, on average, the face value of a firm's debt increases over time suggests that the Black-Cox model as we implement it is misspecified, because the model assumes that default boundary is constant. Despite this the model succeeds in capturing the term structure of default rates to a reasonable degree as Figure 3 shows and the misspecification is therefore "mild" and not crucial for our main result that the model matches investment-grade spreads.

4. Conclusion

Much of the existing literature on testing structural models relies heavily on estimates of default probabilities obtained from historical default frequencies. A much used approach takes the historical default rate for a single rating and maturity as an estimate of the default probability when calculating the spread at that same maturity and rating. We find that the outcome of this approach depends strongly on the historical period from which the default rate is obtained and we show in simulations that a single historical default rate is a very noisy estimator of the default probability. Furthermore, the distribution of the historical default rate for any investment-grade rating is skewed, meaning that the observed historical default rate is likely to be below the ex ante default probability. This in turn implies that when testing a structural model that is calibrated to the historical default rate, one would find predictions of the spread are also likely to appear too low relative to the actual spread, even if the structural model is the true model.

We propose a new method to calibrate structural models to historical default rates. In this approach we extract the default boundary—the value of the firm, measured as a fraction of the face value of debt, below which the firm defaults—by minimizing the difference between actual and model-implied default rates across a wide range of maturities and ratings. We show that this approach dramatically improves the statistical properties of estimated investment-grade default probabilities, in terms of both standard deviation and skewness.

Using our proposed approach we test the Black-Cox model using U.S. data on spreads from individual firms over the period 1987–2012. We find that model spreads match average actual investment-grade credit spreads well. In other words we do not find evidence of a "credit spread puzzle." Going beyond testing the puzzle, we find that the time series of model-implied investment-grade spreads tracks average actual investment-grade spreads well with a correlation of 93%. In contrast, we find that the model significantly underpredicts speculative-grade spreads.

We explore the potential effect of bond illiquidity by sorting pricing errors—the difference between model-implied and actual spreads—on bond age, a proxy for bond illiquidity. We find no relation between pricing errors and bond illiquidity in investment-grade bonds. However, there is a strongly monotone relation between average pricing errors and bond illiquidity in speculative-grade bonds, suggesting that the model underprediction for speculative-grade bonds is due to an illiquidity premium.

Our results show that the credit spread puzzle—the perceived failure of structural models to explain levels of credit spreads for investment-grade bonds—has less to do with deficiencies in the models than with the way in which the models have been implemented. We focus our attention on the Black-Cox model, but our results have important implications for structural models in general.

The results in Huang and Huang (2012) show that many structural models that appear very different in fact generate similar spreads once the models are calibrated to the same historical default rates, recovery rates, and the equity premium. The models tested in Huang and Huang (2012) include features such as stochastic interest rates, endogenous default, stationary leverage ratios, strategic default, time-varying asset risk premiums, and jumps in the firm value process, yet all generate a similar level of credit spread. Although our method of benchmarking historical default rates is different from that of Huang and Huang (2012), we conjecture that, if benchmarked in the way described in this paper, the majority of models considered by Huang and Huang (2012)—at least those with a constant default boundary—would generate spreads similar to each other and, in particular, to those produced by the Black-Cox model.

Appendix A. The Black-Cox Model

We assume that a firm's asset value follows a geometric Brownian motion under the natural measure

$$\frac{dV_t}{V_t} = (\mu - \delta)dt + \sigma dW_t^P, \tag{A1}$$

where δ is the payout rate to debt and equity holders, μ is the expected return on the firm's assets, and σ is the volatility of returns on the asset.

The firm is financed by equity and a single zero-coupon bond with face value F and maturity T. The firm defaults the first time the asset value is below some fraction d of the face value of debt. One interpretation of the default boundary is that the bond has covenants in place that allow bondholders to take over the firm if firm value falls below the threshold. The cumulative default probability in the Black-Cox model at time T is

$$\pi^P(dL,\Theta^P,T)=N\left[-\Big(\frac{(-\log(dL)+\big(\mu-\delta-\frac{\sigma^2}{2}\big)T}{\sigma\sqrt{T}}\Big)\right] + \exp\left(\frac{2\log(dL)(\mu-\delta-\frac{\sigma^2}{2})}{\sigma^2}\right)N\left[\frac{\log(dL)+\big(\mu-\delta-\frac{\sigma^2}{2}\big)T}{\sigma\sqrt{T}}\right], \tag{A2}$$

where $L=\frac{F}{V_0}$ is the leverage and $\Theta^P=(\mu,\sigma,\delta)$ (see Bao 2009). The risk-neutral default probability, π^Q, is obtained by replacing μ with r in Equation (A2).

Appendix B. Firm Data

To compute bond prices in the Merton model, we need the issuing firm's leverage ratio, payout ratio, and asset volatility. This appendix gives details on how we calculate these quantities using CRSP/Compustat.

Firm variables are collected in CRSP and Compustat. To do so, we match a bond's CUSIP with CRSP's CUSIP. In theory the first 6 digits of the bond's CUSIP plus the digits "10" correspond to CRSP's CUSIP, but, in practice, only a small fraction of firms is matched this way. Even if there is a match we check if the issuing firm has experienced merger and acquisition (M&A) activity during the life of the bond. If there is no match, we hand-match a bond's CUSIP with firm variables in CRSP/Compustat.

Leverage ratio: Equity value is calculated on a daily basis by multiplying the number of shares outstanding with the price of shares. Debt value is calculated in Compustat as the latest quarter observation of long-term debt (DLTTQ) plus debt in current liabilities (DLCQ). Leverage ratio is calculated as $\frac{\text{Debt value}}{\text{Debt value+Equity value}}$.

Payout ratio: The total outflow to stake holders in the firm is interest payments to debt holders, dividend payments to equity holders, and net stock repurchases. Interest payments to debt holders is calculated as the previous year's total interest payments (previous fourth quarter's INTPNY). Dividend payments to equity holders is the indicated annual dividend (DVI) multiplied by the number of shares. The indicated annual dividend is updated on a daily basis and is adjusted for stock splits, etc. Net stock repurchase is the previous year's total repurchase of common and preferred stock (previous fourth quarter's PRSTKCY). The payout ratio is the total outflow to stake holders divided by firm value, where firm value is equity value plus debt value. If the payout ratio is larger than 0.13, three times the median payout in the sample, we set it to 0.13.

Equity volatility: We calculate the standard deviation of daily returns (RET in CRSP) in the past 3 years to estimate daily volatility. We multiply the daily standard deviation with $\sqrt{255}$ to calculate annualized equity volatility. If there are no return observations on more than half the days in the 3-year historical window, we do not calculate equity volatility and discard any bond transactions on that day.

Appendix C. Convexity Bias When Using a "Representative Firm" to Calculate Default Probabilities

Our finding in Section 3.3 that the Black-Cox model matches default probabilities for BBB-rated firms including horizons as short as 1 year is surprising, since it is an established stylized fact in the literature that short-run default probabilities in structural models with only diffusion risk are much too low. Papers showing that default probabilities at short horizons are too low include Zhou (2001), Leland (2004), Leland (2006), Cremers, Driessen, and Maenhout (2008), Zhang, Zhou, and Zhu (2009), and McQuade (2013), among others.

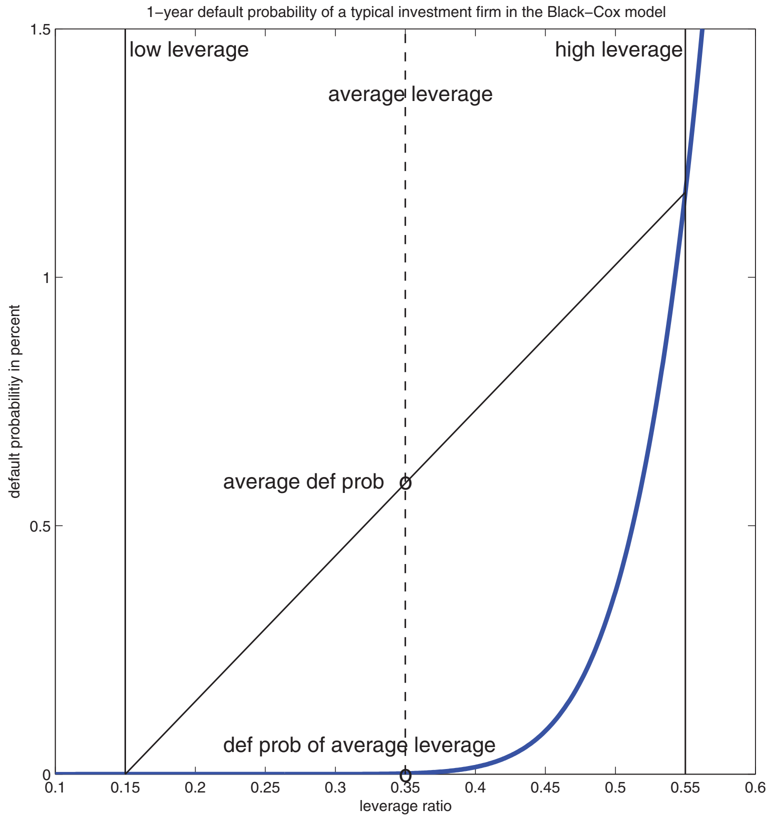

Figure A1
Convexity bias when calculating the default probability in the Black-Cox model using average leverage and comparing it to the average default probability
It is common in the literature to compare historical default rates to model-implied default probabilities, where the latter are calculated using average firm variables. This introduces a bias because the default probability in structural models is a nonlinear function of firm variables. The figure illustrates the bias in case of two firm observations with the same rating, one with a low leverage ratio and one with a high leverage ratio. The two observations can be two different firms at the same point in time or the same firm at two different points in time. Asset volatility is 5%, dividend yield 3.7%, Sharpe ratio 0.22, and risk-free rate 5%.

We arrive at a different conclusion because we allow for cross-sectional variation in asset volatility and both cross-sectional and time-series variation in leverage and payout rates. In contrast, the existing literature uses a "representative firm" with average asset volatility, leverage, and payout rate within a given rating category. Using a representative firm leads to bias due to Jensen's inequality because the default probability is typically convex in asset volatility and leverage (while it is close to linear in the payout rate). Figure A1 illustrates this convexity bias in the case of leverage. The convexity bias when using a representative firm to calculate spreads is known to the literature, but, importantly, the impact of the convexity bias on the short-run default probabilities has not been recognized in the literature.[14]

[14] Bhamra et al. (2010) present a structural-equilibrium model with macroeconomic risk and simulate default rates over 5 and 10 years and find a substantial effect in allowing for firm heterogeneity. They do not look at default probabilities below 5 years, whereas they are our main focus here.

Table A1
Convexity bias when calculating default probabilities in the Black-Cox model using the representative firm approach

Maturity	1	2	3	4	5	6	7	8	9	10
A. True economy (there is variation in leverage ratios)										
Average default probability	0.13	0.59	1.24	1.98	2.75	3.50	4.23	4.92	5.58	6.20
Asset volatility	25.0	25.0	25.0	25.0	25.0	25.0	25.0	25.0	25.0	25.0
B. Representative firm (Average leverage ratio used)										
default probability	0.00	0.00	0.05	0.20	0.49	0.89	1.37	1.90	2.46	3.03
Asset volatility	25.0	25.0	25.0	25.0	25.0	25.0	25.0	25.0	25.0	25.0
C. Representative firm, average def. prob. at bond maturity is matched										
Default probability	0.13	0.59	1.24	1.98	2.75	3.50	4.23	4.92	5.58	6.20
Implied asset volatility	44.3	37.3	34.3	32.6	31.4	30.5	29.9	29.3	28.9	28.6
D. Representative firm, average def. prob. at 10-year bond maturity is matched										
Default probability	0.00	0.03	0.24	0.74	1.47	2.34	3.30	4.28	5.25	6.20
Implied asset volatility	28.6	28.6	28.6	28.6	28.6	28.6	28.6	28.6	28.6	28.6

It is common in the literature to compare average actual default rates to model-implied default probabilities, where model-implied default probabilities are calculated using average firm variables. This introduces a bias because the default probability and spread in the Merton model is a nonlinear function of firm variables. This table shows the magnitude of this bias. Panel A shows, for maturities between one and 10 years, the average default probability for 100,000 firms that have different leverage ratios but are otherwise identical. Their common asset volatility is 25% and payout rate 3.7%. Their leverage ratios are simulated from a normal distribution with mean 0.28 and standard deviation 0.18 (truncated at zero). The risk-free rate is 5%. Panel B shows the default probability of a representative firm where the average leverage ratio is used. In panel C, for each maturity—one at a time—an asset volatility is computed such that, for a representative firm with a leverage ratio equal to the average leverage ratio, the default probability is equal to the average default probability in the economy (given in the first row in panel A and again in panel C). This is done separately for each maturity. The panel shows the resulting implied asset volatility. Panel D shows the results of a calculation similar to that in panel C, except here the asset volatility used to compute the default probability for each maturity is the value that matches the average *10-year* default probability in the economy.

To document the impact of the convexity bias on short-run default probabilities we focus on heterogeneity in leverage and carry out a simulation of 100,000 firms. For each firm, we use an asset volatility of 23%, a payout rate of 3.3%, and a Sharpe ratio of 0.22. The firms differ only in their leverage ratios and we draw 100,000 values from a normal distribution with mean 0.29 and a standard deviation of 0.18.[15] The chosen values are median values for BBB firms, and the standard deviation of leverage in the simulation is equal to the empirical standard deviation of BBB firms in the sample. Finally, the risk-free rate is 5%. For each firm, we calculate the cumulative default probability for different maturities. Panel A in Table A1 shows the average default probability and the correct asset volatility of 23% that is used for all firms and at all maturities.

Zhou (2001), Leland (2004, 2006), and McQuade (2013) use values of the leverage ratio, payout rate, and asset volatility averaged over time and firms to calculate model-implied default probabilities for a representative firm and then compare these with historical averages. To see the extent of the convexity bias in the Black-Cox model when using their approach, we calculate the term structure of default probabilities in panel B of Table A1 for a representative firm with a leverage ratio equal to the mean in our simulation of 0.29. There is a downward bias in default probabilities relative to the correct values given in panel A, and the bias becomes more pronounced at shorter maturities. For example, the 1-year default probability of the representative firm in panel B is 0.00%, whereas the true average default probability in panel A is 0.42%. The aforementioned

15 If a simulated leverage ratio is negative, we set it to zero. This implies that the average leverage ratio is slightly higher than 0.29, namely 0.2937 in our simulation.

papers compare the default probability of the representative firm with the average historical default rate and since the historical default rate reflects the average default probabilities in panel A, their results for particularly short-maturity default probabilities are strongly biased.

Cremers et al. (2008), Zhang et al. (2009), and Huang and Huang (2012) let a representative firm match historical default rates by backing out asset volatility. To examine how the convexity bias influences the implied asset volatility, we proceed as follows. For a given maturity, we compute the asset volatility that allows the representative firm to match the average default probability in the economy at that given maturity. Panel C shows the implied asset volatilities, and we see two problems with this approach. The first problem is that asset volatility is biased: all firms in the economy have an asset volatility of 23%, and yet the implied asset volatility ranges from 27.0% at the 10-year horizon to 43.3% at the 1-year horizon. The finding that implied asset volatility in the diffusion-type structural models is too high, particularly at shorter horizons, has been seen as a failure of the models, but this example shows that the high implied asset volatility mechanically arises from the use of a representative firm. The second problem is that it is not possible to match the term structure of default probabilities without counterfactually changing the asset volatility maturity-by-maturity.

Cremers, Driessen, and Maenhout (2008) and Zhang, Zhou, and Zhu (2009) use a representative firm to imply out asset volatility by matching long-term default rates and then use this asset volatility to calculate the term structure of default probabilities. We replicate this approach by implying out the asset volatility that makes the representative firm's default probability match the average default probability for the 10-year bond in the economy and then calculate the term structure of default probabilities for this representative firm. The implied asset volatility is 27.0% and the term structures are in panel D. The difference between the implied asset volatility of 27.0% and the true value of 23% reflects a moderate convexity bias at the 10-year horizon, but since the bias becomes more severe at shorter horizons, the strong downward bias in default probabilities reappears as maturity decreases. Thus, the bias in short-term default probabilities persists when using a representative firm and imputing asset volatility by matching a long-term default rate.[16]

In summary, we show that the term structure of default probabilities in the Black-Cox model is downward biased, and more so at short maturities, when using a representative firm. This is likely to be true for any standard structural model: default probabilities are strongly bias downward at short maturities. Existing evidence (showing that default probabilities at short horizons are much too low) in Zhou (2001), Leland (2004, 2006), Cremers, Driessen, and Maenhout (2008), Zhang, Zhou, and Zhu (2009), and McQuade (2013) is subject to this strong bias and therefore not reliable.

References

Acharya, V., Y. Amihud, and S. Bharath. 2013. Liquidity risk of corporate bond returns: Conditional approach. *Journal of Financial Economics* 110:358–86.

Bai, H. 2016. Unemployment and credit risk. Working Paper.

Bao, J. 2009. Structural models of default and the cross-section of corporate bond yield spreads. Working Paper.

Bao, J., J. Pan, and J. Wang. 2011. The illiquidity of corporate bonds. *Journal of Finance* 66:911–46.

---

16 Our results clarify those in Bhamra, Kuehn, and Strebulaev (2010). Within the framework of their structural-equilibrium model, they compare a representative firm with a cross-section of firms and find that the slope of the term structure of default probabilities is flatter for the cross-section of firms. In their experiment, the cross-section of firms have an average default probability that is more than three times as large as the default probability of the representative firm (their Table 3, panels B and C). Since the term structure of default probabilities becomes flatter for a representative firm at the same time as default risk increases, it is not clear if it is cross-sectional variation or the rise in default probability that drives the flattening of the term structure. Since we hold the 10-year default probability fixed in panels A and D, it is clear in our analysis that the flatter term structure is driven by cross-sectional variation in leverage alone.

Bhamra, H. S., L.-A. Kuehn, and I. A. Strebulaev. 2010. The levered equity risk premium and credit spreads: A unified framework. *Review of Financial Studies* 23:645–703.

Black, F., and J. Cox. 1976. Valuing corporate securities: Some effects of bond indenture provisions. *Journal of Finance* 31:351–67.

Booth, L., D. Gounopoulos, and F. Skinner. 2014. The choice between callable and noncallable bonds. *Journal of Financial Research* 37:435–59.

Campello, M., L. Chen, and L. Zhang. 2008. Expected returns, yield spreads, and asset pricing tests. *Review of Financial Studies* 21:1297–338.

Chen, H. 2010. Macroeconomic conditions and the puzzles of credit spreads and capital structure. *Journal of Finance* 65:2171–12.

Chen, H., R. Cui, Z. He, and K. Milbradt. 2017. Quantifying liquidity and default risks of corporate bonds over the business cycle. *Review of Financial Studies* 31:852–97.

Chen, L., P. Collin-Dufresne, and R. S. Goldstein. 2009. On the relation between the credit spread puzzle and the equity premium puzzle. *Review of Financial Studies* 22:3367–409.

Christoffersen, P., and D. D. R. Elkamhi. 2017. Rare disasters, credit, and option market puzzles. *Management Science* 63:1341–64.

Collin-Dufresne, P., R. Goldstein, and S. Martin. 2001. The determinants of credit spread changes. *Journal of Finance* 56:2177–2207.

Correia, M., J. Kang, and S. Richardson. 2014. Asset volatility. Working Paper.

Cremers, M., J. Driessen, and P. Maenhout. 2008. Explaining the level of credit spreads: Option-implied jump risk premia in a firm value model. *Review of Financial Studies* 21:2209–42.

Davydenko, S. A. 2013. When do firms default? A study of the default boundary. Working Paper.

Dick-Nielsen, J., P. Feldhütter, and D. Lando. 2012. Corporate bond liquidity before and after the onset of the subprime crisis. *Journal of Financial Economics* 103:471–92.

Duffee, G. R. 1998. The relation between Treasury yields and corporate bond yield spreads. *Journal of Finance* 53:2225–41.

Eom, Y. H., J. Helwege, and J.-Z. Huang. 2004. Structural models of corporate bond pricing: An empirical analysis. *Review of Financial Studies* 17:499–544.

Ericsson, J., J. Reneby, and H. Wang. 2015. Can structural models price default risk? Evidence from bond and credit derivative markets. *Quarterly Journal of Finance* 5:1550007.

Feldhütter, P. 2012. The same bond at different prices: Identifying search frictions and selling pressures. *Review of Financial Studies* 25:1155–206.

Feldhütter, P., and D. Lando. 2008. Decomposing Swap Spreads. *Journal of Financial Economics* 88:375–405.

Friewald, N., R. Jankowitsch, and M. G. Subrahmanyam. 2012. Illiquidity or credit deterioration: A study of liquidity in the US corporate bond market during financial crises. *Journal of Financial Economics* 105:18–36.

Gomes, J., U. Jermann, and L. Schmid. 2016. Sticky leverage. *American Economic Review* 106:3800–28.

He, Z., and K. Milbradt. 2014. Endogenous liquidity and defaultable debt. *Econometrica* 82:1443–508.

Houweling, P., A. Mentink, and T. Vorst. 2005. Comparing possible proxies of corporate bond liquidity. *Journal of Banking and Finance* 29:1331–58.

Huang, J., and M. Huang. 2012. How much of the corporate-Treasury yield spread is due to credit risk? *Review of Asset Pricing Studies* 2:153–202.

Krishnamurthy, A., and A. Vissing-Jorgensen. 2012. The aggregate demand for Treasury debt. *Journal of Political Economy* 120:233–67.

Leland, H. 1994. Corporate debt value, bond covenants, and optimal capital structure. *Journal of Finance* 49:157–96.

———. 2004. Predictions of default probabilities in structural models of debt. *Journal of Investment Management* 2.

———. 2006. *Structural models of corporate financial choice.* Princeton Lectures in Finance.

Lin, H., J. Wang, and C. Wu. 2011. Liquidity risk and expected corporate bond returns. *Journal of Financial Economics* 99:628–50.

McQuade, T. J. 2013. Stochastic volatility and asset pricing puzzles. Working Paper, Harvard University.

Merton, R. 1974. On the pricing of corporate debt: The risk structure of interest rates. *Journal of Finance* 29:449–70.

Moody's. 2002. *Default and recovery rates of corporate bond issuers - A statistical review of Moody's ratings performance 1970-2001.* Moody's Investors Service.

———. 2013. Annual Default Study: Corporate Default and Recovery Rates, 1920-2012. *Moody's Investors Service*, 1–64.

Nagel, S. 2016. The liquidity premium of near-money assets. *Quarterly Journal of Economics* 131:1927–71.

Schaefer, S., and I. Strebulaev. 2008. Structural models of credit risk are useful: Evidence from hedge ratios on corporate bonds. *Journal of Financial Economics* 90:1–19.

Zhang, B. Y., H. Zhou, and H. Zhu. 2009. Explaining credit default swap spreads with the equity volatility and jump risks of individual firms. *Review of Financial Studies* 22:5101–31.

Zhou, C. 2001. The term structure of credit spreads with jump risk. *Journal of Banking and Finance* 25:2015–40.