

Measuring Default Risk Premia from Default Swap Rates and EDFs¹

Antje Berndt,² Rohan Douglas,³ Darrell Duffie,⁴
Mark Ferguson,⁵ and David Schranz⁶

Current Version: July 2008

¹We thank CIBC for default swap data; Moody's KMV, Ashish Das, Jim Herrity, Roger Stein, and Jeff Bohn for access to Moody's KMV EDF data; and Moody's Investor Services for a research grant to Antje Berndt that partially supported her work. We thank seminar participants at Cornell University, The University of Chicago, CMU, UC Santa Barbara, University of Texas at Austin, Fannie Mae, University of Arizona, Penn State University, The Federal Reserve Board of Governors, the Federal Reserve Banks of Atlanta and San Francisco, The University of Florida, and at the 2004 NBER Asset Pricing Summer Conference, 2004 SED meeting, 2004 Bachelier World Congress, 2004 EFA meeting, the Moodys-LBS Credit Risk Conference, the Q Group, and the 2005 ESWC meeting. We are also grateful to Yacine Aït-Sahalia, Gene Fama and Mitchell Petersen for useful discussions, to Jeremy Graveline, Gustavo Manso, Sergey Lobanov, Leandro Saita, and Wei Yang for research assistance, and to Linda Bethel and Sandra Berg for technical assistance.

²Tepper School of Business, Carnegie Mellon University.

³Quantifi Incorporated, New York, New York.

⁴Graduate School of Business, Stanford University.

⁵Quantifi Incorporated, New York, New York.

⁶CIBC, Toronto.

Abstract

This paper estimates the degree of variation over time in the price for bearing exposure to U.S. corporate default risk during 2000-2004, based on the relationship between default probabilities, as estimated by Moody's KMV EDFs, and default swap (CDS) market rates. The default-swap data, obtained through CIBC from 39 banks and specialty dealers, allow us to establish a strong link between actual and risk-neutral default probabilities in the three sectors that we analyze: broadcasting and entertainment, healthcare, and oil and gas. We find dramatic variation over time in risk premia, from peaks in the third quarter of 2002, dropping by roughly 50% to late 2003.

This paper estimates the degree of variation over time in the price for bearing default risk, above and beyond expected default loss, for U.S. corporate debt. We exploit the close but time-varying relationship between default probabilities, as estimated by the Moody's KMV EDF measure, and default swap (CDS) market rates.

Our default-swap data, obtained by CIBC from 39 banks and specialty dealers, allow us to establish a strong link between actual and risk-neutral default probabilities for the 93 firms in the three sectors that we analyzed: broadcasting and entertainment, healthcare, and oil and gas. Based on over 180,000 CDS rate quotes, 5-year EDFs explain over 74% of the variation in 5-year CDS rates across issuers and time, controlling for non-linearity and for sectoral and time fixed effects. For a given default probability, we find substantial variation over time in credit spreads. For example, after peaking in the third quarter of 2002, credit risk premia declined steadily and dramatically through late 2003, when, for a given default probability, credit spreads were on average roughly 50% lower than at their peak. For example, fixing a default probability, CDS rates were 41% lower in December 2003 than in August 2002 in the oil-and gas-sector, 69% lower in the broadcasting-and-entertainment sector, and 49% lower in the healthcare sector.

A potential explanation for these declines in default risk premia is that by mid-2002 the corporate debt market had experienced a reduction in risk-bearing capacity, relative to the amount of risk to be borne, driving risk premia to relatively high levels at that time. This may have been due in part to large default losses in prior months and increases in market volatility, perhaps exacerbated by frictions in the entry of new risk capital. Under this hypothesis, fresh capital flowed into this market over the subsequent months in order to take advantage of the high risk premiums offered, eventually (but not immediately) driving these risk premia down. This is similar to the explanation offered by Froot and O'Connell (1999) for dramatic increases in catastrophic risk insurance premia after major losses of capital, with subsequent slow

declines in premia over time as new capital is attracted into the sector. Consistent with this interpretation, we find that credit risk premia are strongly dependent on general stock-market volatility, as measured by the VIX, after controlling for the influence of firm-specific volatility on default probabilities. This may simply reflect the fact that credit spreads match the supply and demand for risk bearing: when the amount of available capital for bearing default risk is small relative to the level of risk, the price for bearing a given amount of default risk is higher. We are not aware, however, of evidence that the market price of risk in equity markets varies to such a large degree over similarly short periods of time, including this particular period. We will discuss whether the measured reductions in default risk premia could also be influenced by errors in estimating default probabilities, by “reaching for yield” by money managers, or by changes in expected recoveries in the event of default, among other potential explanations.

Our study is based on an extensive database of credit default swap (CDS) rates from CIBC, and on Moody’s KMV estimated default frequency (EDF) data. First panel-regression models, and then arbitrage-free term-structure time-series models, are used to estimate how default risk premia vary over time. Previous work such as Driessen (2005) has established evidence that risk-neutral default probabilities are significantly higher than actual default probabilities. This ratio may be viewed as the proportional premium for bearing default risk. For example, if this ratio is 2.0 (for a particular firm, date, and maturity), then market-based insurance of default would be priced at roughly twice the expected discounted default loss. In terms of the average over time of default risk premia, our evidence is generally consistent with previous work. Our main focus is the degree to which these premia vary over time. We explore various alternative interpretations of our finding that the risk premia do indeed vary dramatically.

A weakness of our study is the lack of data bearing on risk-neutral mean loss given

default (LGD). The highest annual cross-sectional sample mean of loss given default during our sample period was reported by Altman, Brady, Resti, and Sironi (2003) to be approximately 75%. Using 75% as a rough estimate for risk-neutral mean loss given default, our measured relationship between CDS and EDF implies that short-term risk-neutral default probabilities are roughly double their actual-probability counterparts, on average, although this premium is higher for high-quality firms than for low-quality firms, and higher for firms in the broadcasting-and-entertainment sector than for oil-and-gas or healthcare firms. In particular, this ratio was dramatically higher in mid-2002 than in late 2003. If the risk-neutral mean LGD were constant over time, at any particular level, then our results on relative changes over time in default risk premia would be largely unaffected by the assumed level of risk-neutral mean LGD. The results of Altman, Brady, Resti, and Sironi (2003), however, indicate that average realized LGDs tend to be positively correlated with aggregate default rates. As a robustness check, we provide some indication of the potential impact of such correlation on estimated CDS rates.

As an illustrative example, Figure 1, which shows estimated actual and risk-neutral 1-year default probabilities for Disney, is consistent with the typical pattern in our sample of high default risk premia in the third quarter of 2002, particularly in the broadcasting-and-entertainment sector. More generally, Figure 2 shows the median, within the broadcasting-and-entertainment sector, of the estimated ratios of risk-neutral to actual default probabilities at each of three maturities: instantaneously short, one year, and five years.

The remainder of the paper is structured as follows. Section I places our work in the context of prior empirical research on default risk premia. Section II describes our data, including a discussion of the terms of default swap contracts and an overview of the construction of the Moody's KMV EDF measure of default probability. Section III presents panel-regression evidence of a strong relationship between CDS rates

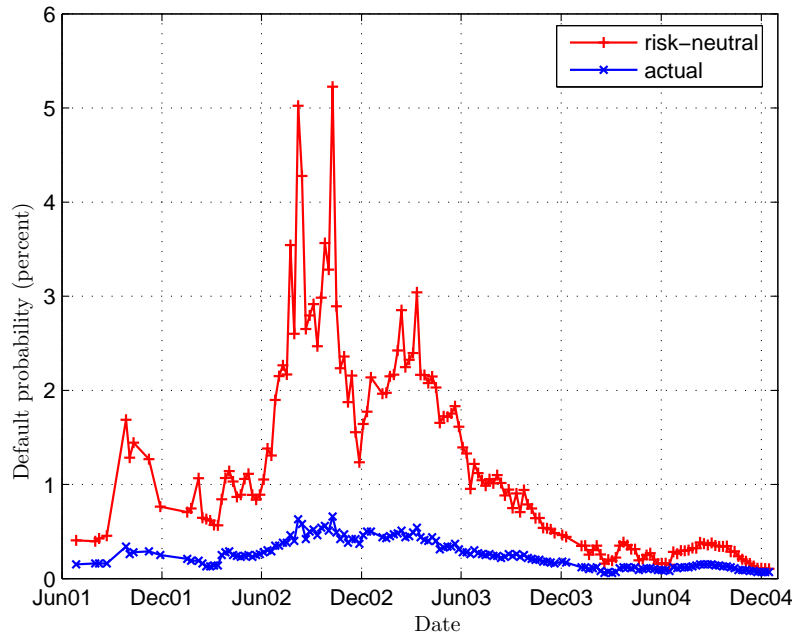


Figure 1: Estimated actual and risk-neutral 1-year default probabilities for Disney.

and EDFs across several sectors, with higher risk premia for high-quality firms, and dramatically declining risk premia from mid-2002 to late 2003. Section IV introduces a time-series model of actual default intensities, and our methodology for maximum-likelihood parameter estimation. Section IV also contains parameter estimates for each firm, based on 12 years of monthly observations of 1-year EDFs for each firm. Section V provides a reduced-form pricing model for default swaps, based on time-series models of actual and risk-neutral default intensities. Section B introduces our parameterization of the time-series model for risk-neutral default intensities, using both EDFs and CDS rates. Section C provides estimates of the parameters for each of the three sectors. Section VI provides a discussion of the implications of the results.

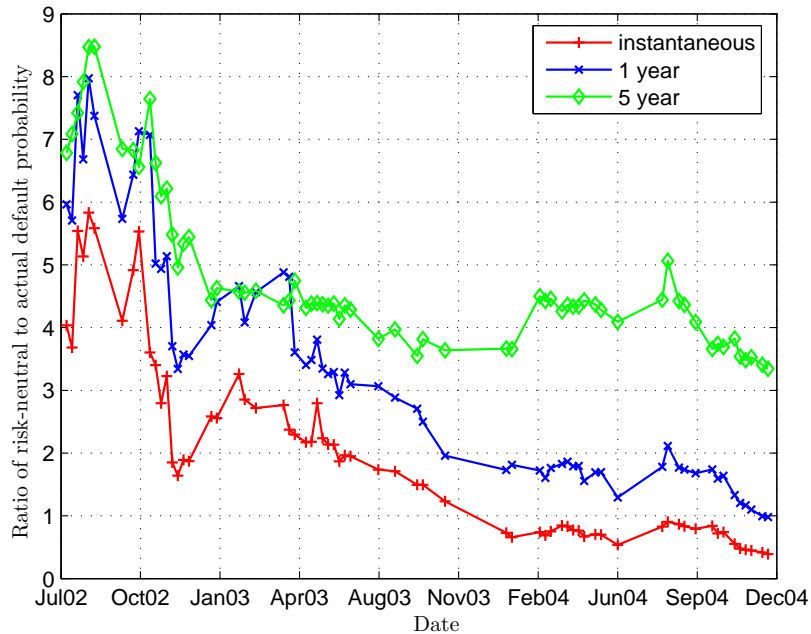


Figure 2: Within-sector medians of the ratios of risk-neutral to actual default probabilities, for the broadcasting-and-entertainment sector, at various maturities.

I. Prior Literature

While Fisher (1959) took a simple regression approach to explaining yield spreads on corporate debt in terms of various credit-quality and liquidity related variables, Fons (1987) gave the earliest empirical analysis, to our knowledge, of the relationship between actual and risk-neutral default probabilities. Driessen (2005) estimated the relationship between actual and risk-neutral default probabilities, using U.S. corporate bond price data (rather than CDS data), and assuming that conditional default probabilities are equal to average historical default frequencies by credit rating. Driessen reported an average ratio of risk-neutral to actual default intensities of 1.89, after accounting for tax and liquidity effects, roughly in line with our estimates. While the conceptual foundations of Driessen’s study are similar to ours, there are substantial differences in our respective research objectives, data sources, and methodology. First, the time periods covered are different. Second, the corporate bonds underlying

Driessen's study are less homogeneous with respect to their sectors, and have significant heterogeneity with respect to maturity, coupon, and time period. Each of our CDS rate observations, on the other hand, is effectively a new 5-year par-coupon credit spread on the underlying firm that is not as corrupted, we believe, by tax and liquidity effects, as are corporate bond spreads. (Driessen estimated the portion of the bond yield spread that is associated with taxes.) Third, and most importantly in terms of our objective of examining time variation in default risk premia, we do not assume that current conditional default probabilities are equal to historical average default frequencies by credit rating. Kavvathas (2001) and others have shown that, for a given firm at a given time, the historical default frequency by firms of the same rating is a stale and coarse-grained estimator of conditional default probability. Moody's KMV EDF measures of default probability provide significantly more power to discriminate among the default probabilities of firms (Kealhofer (2003), Kurbat and Korbalev (2002), Bohn, Arora, and Korablev (2005)). This is not a critical issue for Driessen's results, because he does not focus on time variation in risk premia, but rather reports the average over time of risk premia, by rating.

Blanco, Brennan, and Marsh (2005) show that CDS rates represent somewhat fresher price information than do bond yield spreads. This may be due to the fact that default swaps are "un-funded exposures," in the language of dealers, meaning that in order to execute a trade, neither cash nor the underlying bonds need to be immediately sourced and exchanged. Default swap rates are therefore less likely to be affected by market illiquidity than are bond yield spreads. The extent of this difference in liquidity is explored in Longstaff, Mithal, and Neis (2005). The notional amount of debt covered by default swaps has been roughly doubling each year for the past decade, and in 2007 is estimated to be over 30 trillion U.S. dollars, according to the British Bankers Association (www.bba.org).

Bohn (2000), Delianedis and Geske (1998), Delianedis, Geske, and Corzo (1998),

and Huang and Huang (2003) use structural approaches to estimating the relationship between actual and risk-neutral default probabilities, generally assuming that the Black-Scholes-Merton model applies to the asset value process, and assuming constant volatility. Eom, Helwege, and Huang (2004) have found that these structural models tend to fit the data rather poorly, and typically underestimate credit spreads, especially for shorter maturity bonds. Chen, Collin-Dufresne, and Goldstein (2005) show an improvement in fit by incorporating an assumption of counter-cyclical default boundaries. New work by Saita (2005) estimates the high levels of risk premia that can be obtained for portfolios of corporate debt through diversification.

II. The EDF and CDS Data

This section discusses our data sources for conditional default probabilities and for default swap rates.

A. The EDF Data

Moody's KMV provides its customers with, among other data, current firm-by-firm estimates of conditional probabilities of default over time horizons that include the benchmark horizons of 1 and 5 years. For a given firm and time horizon, this "EDF" estimate of default probability is fit non-parametrically from the historical default frequency of other firms that had the same estimated "distance to default" as the target firm. The distance to default of a given firm is a leverage measure adjusted for current market asset volatility. Roughly speaking, distance to default is the number of standard deviations of annual asset growth by which the firm's expected assets at a given maturity exceed a measure of book liabilities. The liability measure is, in the current implementation of the EDF model, the firm's short-term book liabilities plus one half of its long-term book liabilities. Estimates of current assets and the current standard

deviation of asset growth (“volatility”) are calibrated from historical observations of the firm’s equity-market capitalization and of the liability measure. The calibration, explained for example in Vassalou and Xing (2004), is based on the model of Black and Scholes (1973) and Merton (1974), by which the price of a firm’s equity may be viewed as the price of an option on assets struck at the level of liabilities. Crosbie and Bohn (2002) and Kealhofer (2003) provide more details on the KMV model and the fitting procedures for distance to default and EDF. Bharath and Shumway () show that the fitting procedure is relatively robust. Duffie, Saita, and Wang (2007) and Bharath and Shumway () show that, although distance to default (DD) is a sufficient explanatory variable for conditional default probabilities in the theoretical models of Black and Scholes (1973), Merton (1974), Fischer, Heinkel, and Zechner (1989), and Leland and Toft (1996), among others, some incremental predictive power can be obtained by including some additional firm-specific and macro-economic explanatory variables.

While one could criticize the EDF measure as an estimator of the “true” conditional default probability, it has some important merits for business practice and for our study, relative to other available approaches to estimating conditional default probabilities. First, the Moody’s KMV EDF is readily available for essentially all public U.S. companies, and for a large fraction of foreign public firms. (There is a private-firm EDF model, which we do not rely on, since our CDS data are for public firms.) The EDF is fitted non-parametrically to the distance to default, and is therefore not especially sensitive, at least on average, to model mis-specification. While the measured distance to default is itself based on a theoretical option-pricing model, the function that maps DD to EDF is consistently estimated in a stationary setting, even if the underlying theoretical relationship between DD and default probability does not apply. That is, conditional on only the distance to default, the measured EDF is equal to the “true” DD-conditional default probability as the number of ob-

servations goes to infinity, under typical mixing and other technical conditions for non-parametric qualitative-response estimation.

A common industry measure of default likelihood is the average historical default frequency of firms with the same credit rating as the target firm. This measure is often used, for example, in implementations of the CreditMetrics approach (www.creditmetrics.com), and is convenient given the usual practice by financial-services firms of tracking credit quality by internal credit ratings based on the approach of the major recognized rating agencies such as Moody's and Standard and Poors. The ratings agencies, however, do not claim that their ratings are intended to be a measure of default probability, and they acknowledge a tendency to adjust ratings only gradually to new information, a tendency strongly apparent in the empirical analysis of Behar and Nagpal (1999), Lando and Skødeberg (2002), Kavvathas (2001), and Nickell, Perraudin, and Varotto (2000), among others. This tendency to adjust ratings gradually is illustrated in Figure 3, which shows dramatic variation in default rates by rating depending on whether the prior rating was higher or lower. Bohn, Arora, and Korablev (2005) report that the Moody's KMV EDF measure has an out-of-sample accuracy ratio¹ for our sample period, 2000-2004, of 0.84, as opposed to an accuracy ratio of 0.72 for ratings-based default prediction. Duffie, Saita, and Wang (2007) describe a more elaborate default prediction model, using distance to default as well as other covariates, that achieves an accuracy ratio that is slightly higher than that of the EDF during this period. In Section VI, we discuss the degree to which our results on time variation in default risk premia may be influenced by the accuracy of the EDF measure.

The Moody's KMV EDF measure is also extensively used in the financial services industry. For example, from information provided to us by Moody's KMV, 40 of the world's 50 largest financial institutions are subscribers. Indeed, Moody's KMV is the most widely used name-specific major source of conditional default probability

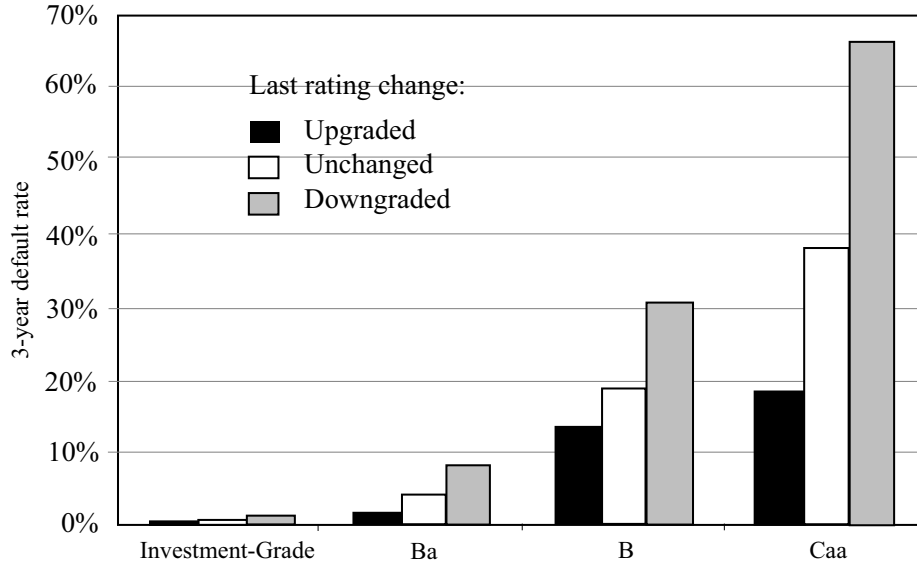


Figure 3: Three-year average default rates by rating, 1996-2003 data. Within rating, average default rates are also shown for the subset of firms in that rating whose prior rating was higher, and the subset whose prior rating was lower. Source: Moody’s, 2004.

estimates of which we are aware, covering over 26,000 publicly traded firms.

Our basic analysis in Section 3 directly relates daily observations of 5-year CDS rates to the associated daily 5-year EDF observations. For our time-series model of default intensities, however, we turn in Section IV to 12 years of monthly observations of 1-year EDFs. By sampling monthly rather than daily, we mitigate equity market microstructure noise, including intra-week seasonality in equity prices, and we also avoid some of the intra-month seasonality in EDFs caused by monthly uploads of firm-level accounting liability data. By using 1-year EDFs rather than 5-year EDFs, our intensity estimates are less sensitive to model mis-specification, as the 1-year EDF is theoretically much closer to the intensity than is the 5-year EDF. As a robustness check, we have also fit our time-series model of default risk premia to 5-year EDF data; the results are similar in major respects to those reported here.

B. Default Swaps and the CDS Database

A default swap, often called (with inexplicable redundancy) a “credit default swap” (CDS), is an over-the-counter derivative security designed to transfer credit risk.

With minor exceptions, a default swap is economically equivalent to a bond insurance contract. The buyer of protection pays periodic (usually quarterly) insurance premiums until the expiration of the contract or until a contractually defined credit event, whichever is earlier. For our data, the stipulated credit event is default by the named firm. If the credit event occurs before the expiration of the default swap, the buyer of protection receives from the seller of protection the difference between the face value and the market value of the underlying debt. The buyer of protection normally has the option to substitute other types of debt of the underlying named obligor. The most popular settlement mechanism at default is for the buyer of protection to submit to the seller of protection debt instruments of the named firm, of the total notional amount specified in the default-swap contract, and to receive in return a cash payment equal to that notional amount, less the fraction of the default-swap premium that has accrued (on a time-proportional basis) since the last regular premium payment date. Recently, the market has introduced an auction for cash settlement of CDS for cases involving major defaults, such as those in 2005 of Collins-Aikman and Delphi, in order to avoid settlement disruptions caused by a shortage of transferable debt instruments of the underlying name, relative to the number and sizes of required settlement trades.

The CDS rate is the annualized premium rate, as a fraction of the face value of debt covered. Using an actual-360 day-count convention, the CDS rate is thus four times the quarterly premium. Our observations are at-market, meaning that they are bids or offers of the default-swap rates at which a buyer or seller of protection is proposing to enter into new default swap contracts, without an up-front payment. Because there is no initial exchange of cash flows on a standard default swap, the at-market CDS rate is, in theory, that for which the net market value of the contract is zero. In practice, there are implicit dealer margins that we treat by assuming that the average of the bid and ask CDS rates is the rate at which the market value of the

default swap is indeed zero.

For the purpose of settlement of default swaps, the contractual definition of default normally allows for bankruptcy, a material failure by the obligor to make payments on its debt, or a restructuring of its debt that is materially adverse to the interests of creditors. (This is the same definition of default used for purposes of the Moody's KMV EDF estimator of default probability.) The coverage of default swaps for restructuring has been a question of debate among the community of buyers and sellers of protection. Banks, especially European banks, generally prefer to include restructuring as a covered default event, given the relatively greater exposure of bank loans (versus traded bonds) to restructuring risk. ISDA, the industry coordinator of standardized default-swap contracts, has arranged a consensus contractual definition of default and coverage in the event of default that is likely to be reflected in most of our data. This consensus definition has been adjusted over time, however, and to the extent that these adjustments during our observation period are material, or to the degree of heterogeneity in our data over the definition of default that is applied, our results could be affected. The contractual definition of default can affect the estimated risk-neutral implied default probabilities, since of course a wider definition of default implies a higher risk-neutral default probability.

For a given level of seniority (our data are based on senior unsecured debt instruments), there is less recovery-value heterogeneity if the event of default is bankruptcy or failure to pay, for these events normally trigger cross-acceleration covenants that cause debt of equal seniority to convert to immediate obligations that are *pari passu*, that is, of equal legal priority. If restructuring is included as a contractually covered credit event, however, then there is the potential for significant heterogeneity at default in the market values of the various debt instruments of the obligor, as fractions of their respective principals, especially when there is significant heterogeneity with respect to maturity. The resulting cheapest-to-deliver option can therefore increase

the loss to the seller of protection in the event of default. Without, at this stage, data bearing on the heterogeneity of market value of the pool of deliverable obligations for each default swap, we are in effect treating the cheapest-to-deliver option value as a constant that is absorbed into the estimated risk-neutral mean fractional loss given default, L^* , to the seller of protection in the event of default. While we vary L^* across sectors, we generally assume that L^* is constant across time. To the extent that L^* varies over time or across issuers, our implied risk-neutral default probabilities would be corrupted.

The impact of the cheapest-to-deliver option is, within the current “modified” and “modified-modified” ISDA contractual standards, mitigated by a contractual restriction on the types of deliverable debt instruments, especially with respect to maturity. While there is a tendency for a different standard for European (usually “modified-modified”) versus U.S. firms (usually “modified”), all of our data are for U.S. firms. Restructurings are also associated with higher average default recoveries.²

Ignoring the cheapest-to-deliver effect, the CDS rate is, in frictionless markets, a close approximation of the par-coupon credit spread of the same maturity as the default swap, due to arbitrage reasoning shown by Duffie (1999). Our results thus speak to the relationship between EDFs and corporate credit spreads. Indeed, that par credit spreads are relatively close to CDS rates is confirmed in the empirical analysis of Blanco, Brennan, and Marsh (2005), provided one measures bond spreads relative to interest-rate swap yields, rather than treasury yields, which can be contaminated by tax exemption of coupon income, repo specials, and liquidity effects. To the extent that CDS rates differ from bond credit spreads, Blanco, Brennan, and Marsh (2005) indicate that CDS rates tend to reflect slightly fresher information.

Our CIBC data set consists of over 180,000 intra-day CDS rate quotes on 93 firms from three Moody’s industry groups. The anonymous sources of these quotes include 27 investment banks and 12 default-swap brokers. The concentration of the number

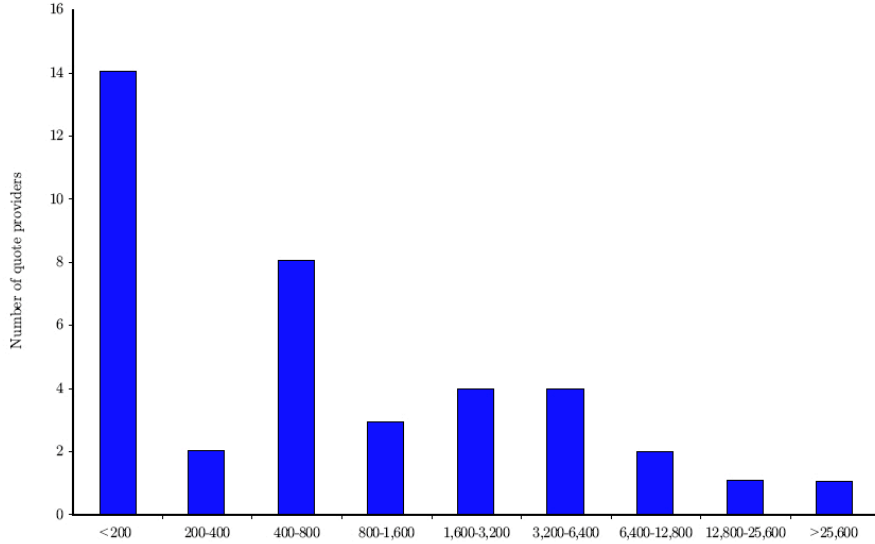


Figure 4: Distribution of CDS quote providers by number of quotes provided. Data source: CIBC.

of quotes by source is shown in Figure 4. A weakness of our study is that a single broker-dealer is the source for almost 60% of our observations.

We selected three representative Moody's-defined North American industry groups: Broadcasting and Entertainment, Oil and Gas, and Healthcare. The CDS quotes are for 1-year and 5-year, quarterly premium, senior unsecured, US-Dollar-denominated, at-the-money default swaps. The 5-year quotes are the most liquid, and are the basis for our panel-regression analysis. (Our main time-series analysis relies more heavily on the 5-year benchmark, allowing for noisy observation of the 1-year quotes. We do explore a variant, described in Section 4, in which both the 1-year and the 5-year CDS are assumed to be measured with error.) A company from any of these three sectors is included in our time-series analysis if and only if at least 1,000 historical pairs of CDS bid and ask quotes for that firm were available during the sample period. The range of credit qualities of the included firms may be judged from Figure 5, which shows, for each credit rating, the number of firms in our study of that median Moody's rating during the sample period. Figure 5 indicates a concentration of Baa-rated firms. Daily CDS mid-point rate quotes were estimated from intra-day bid and ask quotes.³

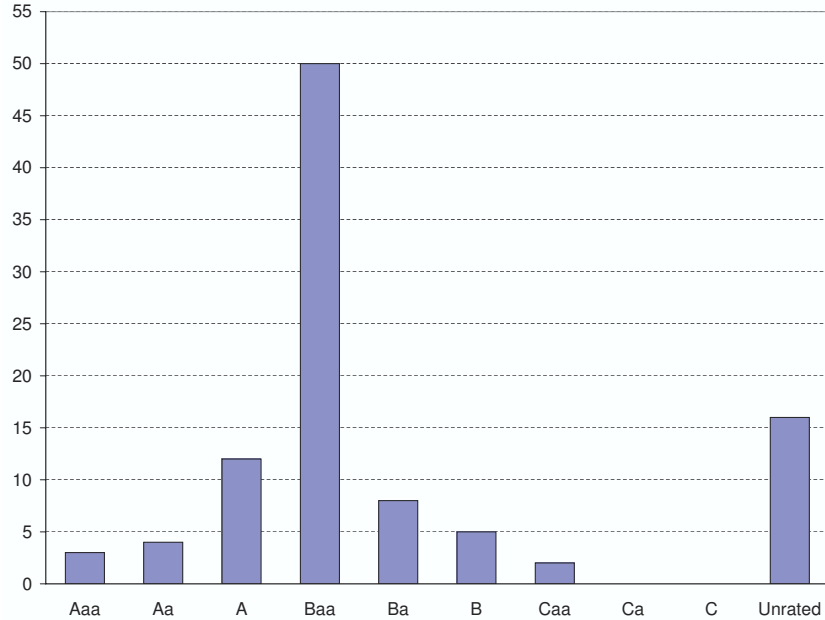


Figure 5: Number of firms by median credit rating during the sample period. Sources: CIBC and Moody's.

The firms that we studied from the broadcasting-and-entertainment industry are listed in Table I, along with their median 1-year EDF and median Moody's credit rating during the sample period from June 2000 to December 2004, and the number of CDS quotes available for each. The same information covering firms from the healthcare and oil-and-gas industries is provided in Appendix B.

III. Panel Regression Analysis

In order to inform the parameterization of our time-series model, and to obtain a simple and relatively robust measure of the sensitivity of credit spreads (CDS rates) to default probabilities, we undertook a panel-regression analysis of all 33,912 paired EDF and median⁴ CDS observations from December 2000 through December 2004, for our 3 sectors. Outliers that could be identified unambiguously were removed manually.

A simple preliminary ordinary-least-squares (OLS) linear model of the relationship

Table I: Broadcasting and Entertainment Firms

Name of Firm	Median EDF (basis points)	Median Rating	No. Quotes
Adelphia Communications Corp	378	N/A	228
Belo Corp	6	Baa3	1,168
Brunswick Corp	9	Baa2	1,390
Charter Communications Inc	600	N/A	456
Clear Channel Communications Inc	41	Baa3	3,330
Comcast Cable Communications	–	Baa3	1,182
Comcast Corp	40	Baa3	2,723
Cox Communications Inc	17	Baa3	4,956
Cox Enterprises Inc	–	Baa3	1,058
Historic TW Inc	–	Baa1	1,462
Interpublic Group of Cos Inc	229	Baa3	1,095
Knight-Ridder Inc	3	A2	1,290
LibertyMediaCorp	48.5	Baa3	2,244
Mediacom Communications Corp	857	Caa1	168
News America Holdings	–	N/A	1,165
News America Inc	–	Baa3	1,679
OmnicomGroup	38	Baa1	2,539
Primedia Inc	939.5	B3	332
Royal Caribbean Cruises Ltd	107	Ba2	1,043
Sabre Holdings Corp	64.5	Baa2	1,467
Time Warner Inc	135	Baa1	5,549
Viacom Inc	18	A3	3,997
Walt Disney Co	23	Baa1	4,459

between a firm's 5-year CDS rate and its annualized 5-year EDF, measured in basis points on the same day, reveals that the CDS rate increases on average by roughly 16 basis points for each 10 basis point increase in the 5-year EDF. If one were to take the risk-neutral expected loss given default to be, say, 75% and the annual conditional default probabilities to be constant over time, this would imply an average ratio of risk-neutral to actual annual default probabilities of approximately $(16/0.75)/10$, or 2.0, roughly consistent with the results of Driessen (2005). The associated coefficient of determination, R^2 , is 0.73.

Linearity of the CDS-EDF relationship, however, is placed in doubt by a sizable intercept estimate of roughly 33 basis points, more than 30 times its standard error. Absent an unexpectedly large liquidity impact on CDS rates, the fitted default swap rate should be closer to zero at low levels of EDF. While there may be mis-specification

due to the assumed homogeneity of the relationship over time and across firms, we have verified with sector and time fixed effects that the associated intercept estimates are unreasonably large in magnitude. Scatter plots of the CDS-EDF relationship also reveal a pronounced concavity at low levels of EDF. That is, the sensitivity of credit spreads to a firm's default probability seems to decline as default probabilities increase. There is also apparent heteroskedasticity, with dramatically greater variance for higher EDFs. The slope of the fit is thus heavily influenced by the CDS-to-EDF relationship for lower-quality firms.

In order to mitigate the effects of non-linearity and heteroskedasticity, we considered the log-log specification⁵

$$\log Y_i = \alpha + \beta \log X_i + z_i, \quad (1)$$

where (Y_i, X_i) is the i -th observed matched pair of 5-year CDS rate and 5-year EDF for the same firm on a given date, for coefficients α and β , and a residual z_i . The fit, which has an R^2 of 0.69, is illustrated in Figure 6, showing much less heteroskedasticity. (One notes granularity associated with integer variation in EDFs of extremely high-quality firms.) One might have considered a model in which the CDS rate is fit to both 5-year and 1-year EDF observations, given the potential for additional influences of near-term default risk on CDS rates. We have found, however, that the 1-year and 5-year EDFs are extremely highly correlated. As might be expected, adding 1-year EDFs to the regression has no major impact on the quality of fitted CDS rates, and involves substantial noise in the slope coefficients.

We also control for changes in the CDS-to-EDF relationship across time and across sectors. Appendix Table X presents the results of a regression of the logarithm of the daily median CDS rate on the logarithm of the associated daily 5-year EDF observation, including dummy variables for sectors and months. (The oil-and-gas

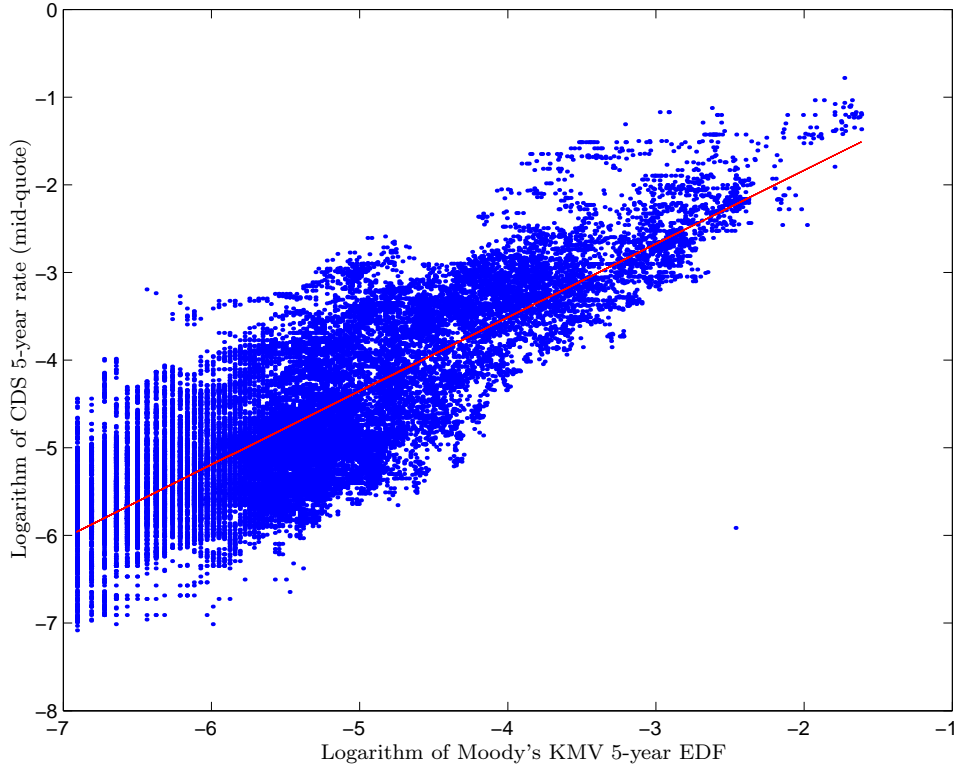


Figure 6: Scatter plot of EDF and CDS observations, logarithmic, and OLS fitted relationship. Source: CIBC (CDS) and Moody's KMV (EDF).

sector for December 2004 is the reference sector and month.) With an R^2 of 74.4%, the fitted model for the oil-and-gas sector may be summarized as

$$\log \text{CDS}_i = 1.45 + 0.76 \log \text{EDF}_i + \sum_{j,k} \hat{\beta}_{j,k} D_{j,k}(i) + z_i, \quad (2)$$

(0.046) (0.015)

where $\hat{\beta}_{j,k}$ denotes the estimate for the dummy multiplier for month j and sector k , tabulated in Appendix Table X, $D_{j,k}(i)$ is 1 if observation i is from month j and sector k and zero otherwise, and z_i denotes the residual. The standard-error estimates reported here and in Appendix Table X are “robust” to heteroskedasticity and correlation of disturbances, using the usual generalized-least-squares estimator for the covariance matrix of regressor coefficients for panel-data regressions, found, for example, in Woolridge (2002), Section 7.8.4.

The standard error for (2) of approximately 0.52, and an assumption of normally distributed disturbances, imply a one-standard-deviation confidence band for a given CDS rate of between 59% and 169% of the fitted rate. While the CDS data are noisy in this sense, the relationship between CDS and EDF is highly significant, and variation in EDF on its own explains a large fraction (an R^2 of about 69%, before controlling for time and sector effects) of variation in CDS rates. For the reference sector and month, oil and gas for December, 2004, five-year mean default rates of 10, 110, and 210 basis points per year are associated⁶ with estimated CDS rates, assuming normality of disturbances, of approximately 28, 161, and 283 basis points, respectively. While the linear-in-logs model captures the apparent declining marginal impact of EDFs on CDS rates, if one were to apply this model to sufficiently high EDFs (above our maximum EDF observation of 2000 basis points), it would eventually imply risk-neutral default probabilities that are below actual default probabilities. Indeed, Figure 6 seems to illustrate a tendency for the linear-in-logs model to understate CDS rates at the highest observed EDFs. A slightly more flexible non-linear model might be preferred.

Figure 7 shows, for each sector k , variation over month j of $e^{\hat{\beta}_{j,k}}$, an estimate of the proportional variation over time of risk premia. That is, $e^{\hat{\beta}_{j,k}}$ is the ratio of the fitted default swap rate for a firm in sector k at month j , to that of an oil-and-gas firm with the same default probability in December 2004. (The maximum of the standard errors of the dummy-variable coefficients shown in Table X is 0.03, indicating a proportional standard deviation in the estimate of $e^{\hat{\beta}_{j,k}}$ of approximately 3% or lower.) Figure 7 indicates dramatic variation over time in risk premia. From August 2002 to December 2003, for a fixed default probability, the estimated reduction⁷ in CDS rates is 41% for the oil-and-gas sector, 69% for the broadcasting-and-entertainment sector, and 49% for the healthcare sector. The broadcasting-and-entertainment sector, in particular, shows dramatic reductions in risk premia from mid 2002 (around the times of default of Adelphia and Worldcom) to late 2003. Section 5 provides additional evidence of

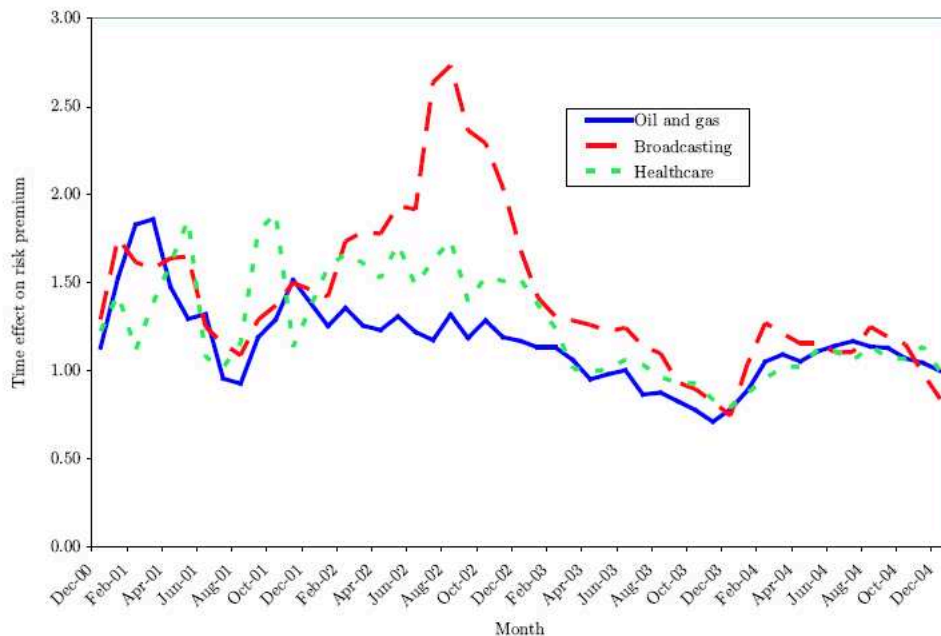


Figure 7: The multipliers for estimated 5-year CDS rates, over time and sector, at a fixed 5-year EDF. These are the exponentials of the dummy coefficients in the log-CDS-to-log-EDF model (2).

variation over time of default risk premia that is revealed through fitted time-series models of actual and risk-neutral default intensities.

Some of the cross-sector differences in CDS rates may due to sectoral differences in the behavior of loss given default. For example, assuming that the ratio of the risk-neutral mean loss given default in the oil-and-gas sector to another sector is the same as the ratio of the empirical average historical loss given default reported by Moody's⁸ for 1982 to 2003, broadcasting-entertainment CDS rates would be approximately $62\%/52\% - 1 = 19\%$ higher than those of the oil-and-gas sector, for equal risk-neutral default probabilities. Similarly, healthcare spreads would be approximately $67\%/52\% - 1 = 29\%$ higher than those of the oil-and-gas sector, for equal risk-neutral default probabilities.

IV. Default Intensity Time-Series Model

We turn to the time-series behavior of actual and risk-neutral default probabilities, measured through firm-level default intensities. We adopt a parameterization that allows a decomposition of time variation in risk premia due to changes in firm-specific default probabilities and due to changes in market-wide default risk premia. In this section, we focus exclusively on actual default intensities. The next section focuses on risk-neutral default intensities.

The default intensity of an obligor is the mean arrival rate of default, conditional on all current information. To be slightly more precise, we suppose that default for a given firm occurs at the first event time of a (non-explosive) counting process N with intensity process λ , relative to a given probability space (Ω, \mathcal{F}, P) and information filtration $\{\mathcal{F}_t : t \geq 0\}$ satisfying the usual conditions. In this case, so long as the obligor survives, we say that its default intensity at time t is λ_t . Under mild technical conditions, this means that, conditional on survival to time t and all information available at time t , the probability of default between times t and $t+h$ is approximately $\lambda_t h$ for small h . We also adopt the relatively standard simplifying doubly-stochastic, or Cox-process, assumption, under which the conditional probability at time t , for a currently surviving obligor, that the obligor survives to some later time T , is

$$p(t, T) = E \left(e^{-\int_t^T \lambda(s) ds} \middle| \mathcal{F}_t \right). \quad (3)$$

For our analysis, we ignore mis-specification of the EDF model itself, by assuming that $1 - p(t, t+1)$ is indeed the current 1-year EDF. From the Moody's KMV data, then, we observe $p(t, t+1)$ at successive dates $t, t+h, t+2h, \dots$, where h is one month. From these observations, we estimate a time-series model of the underlying intensity process λ , for each firm. (Econometrically, this is essentially the same as estimating the time-series behavior of a short-term interest-rate process from one-year

zero-coupon bond prices in an economy with no interest-rate risk premia.)

After some preliminary diagnostic analysis of the EDF data set, we opted to specify a model under which the logarithm $X_t = \log \lambda_t$ of the default intensity satisfies the Ornstein-Uhlenbeck equation

$$dX_t = \kappa(\theta - X_t) dt + \sigma dB_t, \quad (4)$$

where B is a standard Brownian motion, and θ , κ , and σ are constants to be estimated. The behavior for $\lambda = e^X$ is sometimes called a Black-Karasinski model.⁹ This leaves us with a vector $\Theta = (\theta, \kappa, \sigma)$ of unknown parameters to estimate from the available monthly EDF observations of a given firm. We have 144 months of 1-year EDF observations for most of the firms in our sample, for the period January, 1993, to December, 2004.

In general, given the log-autoregressive form (4) of the default intensity, there is no closed-form solution available for the 1-year EDF, $1 - p(t, t + 1)$, from (3).¹⁰ We therefore rely on numerical lattice-based calculations of $p(t, t + 1)$. We have implemented the two-stage procedure for constructing trinomial trees proposed by Hull and White (1994).

The maximum likelihood estimator (MLE) $\hat{\Theta}$ of the parameter vector Θ is then obtained, firm by firm, using a fitting algorithm described in the appendix. That is, for a given firm, $\hat{\Theta}$ solves

$$\sup_{\Theta} \mathcal{L}(\{1 - p(t_i, t_i + 1) : 1 \leq i \leq N\}; 1 - p(t_0, t_0 + 1), \Theta),$$

where t_0, t_1, \dots, t_N are the $N + 1$ observation times for the given firm, and \mathcal{L} denotes the likelihood score of observed EDFs conditioned on the first observation and given Θ . This is not a routine MLE for a discretely-observed Ornstein-Uhlenbeck model, for several reasons:

1. Evaluation of the likelihood score requires a numerical differentiation of the modeled EDF,

$$G(\lambda(t); \Theta) = 1 - E_{\Theta} \left(e^{-\int_t^{t+1} \lambda(s) ds} \mid \lambda(t) \right),$$

where E_{Θ} denotes expectation associated with the parameter vector Θ .

2. As indicated by Kurbat and Korbalev (2002), Moody's KMV caps its 1-year EDF estimate at 20%. Since this truncation, if untreated, would bias our estimator, we explicitly account for this censoring effect on the associated conditional likelihood, as explained in Appendix A.
3. Moody's KMV also truncates the EDF below at 2 basis points. Moreover, there is a significant amount of integer-based granularity in EDF data below approximately 10 basis points, as indicated in Figure 6. We therefore remove from the sample any firm whose sample-mean EDF is below 10 basis points. This leaves us with a sample of 84 firms.
4. There were occasional missing data points. These gaps were also treated exactly, assuming that the event of censoring is independent of the underlying missing observation, as explained in Appendix A.
5. For a small number of firms, an exceptional 1-month fluctuation in the 1-year EDF generated an obviously unrealistic estimate of the mean-reversion parameter κ for that company. We ignored Enron's data point for December 2002, the month it defaulted. Similarly, Magellan Health Services filed for protection under Chapter 11 in March 2003 (we used the EDFs through February 2003), and Adelphia Communications petitioned for reorganization under Chapter 11 in June 2002 (we used the EDFs through May 2002). For Forest Oil, we ignored the outlier months of January and February 1993. Finally, we removed Dynegy

from our data set as its 1-year EDF is capped at 20% for most of 2002 and 2003.

Table VII of Appendix A lists the firms for which we have EDF data, showing the number of monthly observations for each as well as the number of EDF observations that were truncated at 20%. The estimated parameter vector for each firm is provided in Table XI, found in Appendix B. One notes significant dispersion across firms in the estimated parameters. Our Monte-Carlo analysis revealed substantial small-sample bias in the MLE estimators. (See Table XII in Appendix B). We therefore obtain sector-by-sector estimates for κ and σ , while allowing for a firm-specific long-run mean parameter θ . Towards this end, we introduce a joint distribution of EDFs across firms in a given industry sector by imposing joint normality of the Brownian motions driving each firm's EDFs, with a flat cross-firm correlation structure within the sector. In particular, we generalize Equation (4) by assuming that the logarithm $X_t^i = \log \lambda_t^i$ of the default intensity of firm i satisfies the Ornstein-Uhlenbeck equation

$$dX_t^i = \kappa (\theta^i - X_t^i) dt + \sigma \left(\sqrt{\rho} dB_t^c + \sqrt{1 - \rho} dB_t^i \right), \quad (5)$$

where B^c and B^i are independent standard Brownian motions, independent of $\{B^j\}_{j \neq i}$, and the constant pairwise within-sector correlation coefficient ρ is an additional parameter to be estimated. The sector-by-sector estimates of the extended parameter vector

$$\Theta = (\{\theta^i\}, \kappa, \sigma, \rho)$$

obtained from an EM algorithm with Gibbs sampling¹¹ are shown in Table II and in Table XIII in Appendix B. The intensity λ is measured in basis points per year.

Table II: Sector EDF-implied default intensity parameters. Asymptotic standard errors, holding $\theta^i = \hat{\theta}^i$ constant, are in parenthesis.

	$\hat{\kappa}$	$\hat{\sigma}$	$\hat{\rho}$	mean ($\hat{\theta}^i$)	no. firms
Oil and Gas	0.469 (0.045)	1.224 (0.040)	0.282 (0.026)	3.110 –	40
Healthcare	0.470 (0.054)	1.252 (0.043)	0.128 (0.018)	3.028 –	25
Broadcasting & Entertainment	0.427 (0.068)	1.232 (0.054)	0.232 (0.027)	3.591 –	19

V. Risk-Neutral Intensity from CDS and EDF

This section explains our methodology for estimating a joint model of actual and risk-neutral default intensities from CDS and EDF data. The basic idea of the model is that the risk-neutral default intensity of a given firm is a function of its own default intensity, a measure of aggregate default risk in the sector, and a latent variable that captures variation in default risk premia not already captured by the first two variables. Time-series variation in CDS rates, coupled with the behavior of actual default intensities already estimated from the model described in the previous section, is then used to estimate actual and risk-neutral dynamics, and to identify the outcomes of the latent variables.

A. Default Swap Pricing

We begin with a simple reduced-form arbitrage-free pricing model for default swaps. Under the absence of arbitrage and market frictions, and under mild technical conditions, there exists a “risk-neutral” probability measure, also known as an “equivalent martingale” measure, as shown by Harrison and Kreps (1979) and Delbaen and Schachermayer (1999). In our setting, markets should not be assumed to be complete, so the martingale measure is not unique. This pricing approach nevertheless

allows us, under its conditions, to express the price at time t of a security paying some amount, say W , at some bounded stopping time $\tau > t$, as

$$S_t = E^Q \left(e^{-\int_t^\tau r(u) du} W \mid \mathcal{F}_t \right), \quad (6)$$

where r is the short-term interest-rate process¹² and E^Q denotes expectation with respect to an equivalent martingale measure Q , that we fix. One may view (6) as the definition of such a measure Q . The idea is that the actual measure P and the risk-neutral measure Q differ by an adjustment for risk premia.

Under our earlier assumption of default timing according to a default intensity process λ (under the actual probability measure P that generates our data), Artzner and Delbaen (1992) show that there also exists a default intensity process λ^* under Q . Even though we have assumed the doubly-stochastic property under P , this need not imply the same convenient doubly-stochastic property under Q as well. Indeed, Kusuoka (1999) gave a counterexample. We will nevertheless assume the doubly-stochastic property under Q . (Sufficient conditions are given in Duffie (2001), Appendix N.) Thus, we have

$$Q(\tau > T \mid \mathcal{F}_t) = p^*(t, T) = E^Q \left(e^{-\int_t^T \lambda^*(u) du} \mid \mathcal{F}_t \right), \quad (7)$$

provided the firm in question has survived to t .

For convenience, we assume independence, under Q , between interest rates on the one hand, and on the other the default time τ and loss given default. We have verified that, except for levels of volatility of r and λ^* far in excess of those for our sample, the role of risk-neutral correlation between interest rates and default risk is in any case negligible for our parameters. This is not to suggest that the magnitude of the correlation itself is negligible. (See, for example, Duffie (1998).) It follows from (6) and this independence assumption that the price of a zero-coupon defaultable bond

with maturity T and zero recovery at default is given by

$$d(t, T) = \delta(t, T)p^*(t, T), \quad (8)$$

where $\delta(t, T) = E_t^Q \left(e^{-\int_t^T r(s) ds} \right)$ is the default-free market discount factor, and $p^*(t, T)$ is the risk-neutral conditional survival probability of (7). Extensions to the case of correlated interest rates and default times were first treated by Lando (1998).

A default swap stipulates quarterly payments by the buyer of protection at a stipulated annual rate of c , as a fraction of notional, until the default-swap maturity or default, whichever is first. From (8), the market value of the payments by the buyer of protection at the origination date of a default swap of unit notional size is thus $cg(t)$, where

$$g(t) = \frac{1}{4} \sum_{i=1}^n \delta(t, t(i))p^*(t, t(i)), \quad (9)$$

for payment dates $t(1), \dots, t(n)$. The market value of the potential payment by the seller of protection is

$$h(t, c) = E^Q \left(\delta(t, \tau) W_\tau^c 1_{\tau \leq t(n)} \mid \mathcal{F}_t \right), \quad (10)$$

where the payment W_t^c at default, if default occurs at time t , is

$$W_t^c = L_t^* - c \left(t - \frac{\lfloor 4t \rfloor}{4} \right), \quad (11)$$

where $\lfloor x \rfloor$ denotes the largest integer less than x , and where L_t^* denotes the risk-neutral expected fractional loss of notional at time t , assuming immediate default.¹³

The second term in (11) is a deduction for accrued premium.

The current CDS rate is that choice $C(t)$ for the premium rate c at which the market values of the payments by the buyer and seller of protection are equal. That

is, $C(t)$ solves

$$C(t)g(t) = h(t, C(t)). \quad (12)$$

Noting that $h(t, c)$ is linear with respect to c , this is a linear equation to solve for $C(t)$.

We turn to the calculation of $h(t, c)$. By the doubly-stochastic property (see, for example, Duffie (2001), Chapter 11), we first condition on (λ^*, L^*) , and then use the conditional risk-neutral density $e^{-\int_t^s \lambda^*(u) du} \lambda^*(s)$ of τ at time s to get

$$h(t, c) = \int_t^{t^{(n)}} \delta(t, s) E^Q \left(e^{-\int_t^s \lambda^*(u) du} \lambda^*(s) W_s^c \mid \mathcal{F}_t \right) ds. \quad (13)$$

We take L^* to be constant and use, as a numerical approximation of the integral in (13),

$$h(t, c) \simeq \sum_{i=1}^n \delta \left(t, \frac{t^{(i)} + t^{(i-1)}}{2} \right) [p^*(t, t^{(i-1)}) - p^*(t, t^{(i)})] \left(L^* - \frac{c}{8} \right), \quad (14)$$

which involves a time discretization of the integral in (13) that, in effect, approximates between quarter ends with a linear discount function and risk-neutral survival function. Then $C(t)$ is calculated from (12) using this approximation. The discount factors $\delta(t, s)$ are fit from contemporaneous market LIBOR and swap rate data.

B. Model Specification

We now specify a parametric model of the risk-neutral default intensity process $\lambda^{*,i}$ of any given firm i . Motivated in part by the success of our linear-in-logarithms regressions reported in Section 3, we suppose that

$$\log \lambda_t^{*,i} = \beta_0 + \beta_1 \log \lambda_t^i + \beta_2 \log v_t + u_t^i, \quad (15)$$

where β_0, β_1 , and β_2 are constants, $X^i = \log \lambda^i$ is as specified earlier by (5), and v is the geometric average of default intensities $\{\lambda^i\}_{i \in J}$, over a benchmark subset J of large liquid firms in the industry group,¹⁴ in that

$$\log v_t = \frac{1}{|J|} \sum_{i \in J} X_t^i. \quad (16)$$

We suppose that

$$du_t^i = \kappa_u(\theta_u - u_t^i) dt + \sigma_u \sqrt{\rho_u} d\xi_t^c + \sigma_u \sqrt{1 - \rho_u} d\xi_t^i, \quad (17)$$

where θ_u, κ_u , and σ_u are constants, ρ_u is a constant correlation parameter, and where, under the actual probability measure P , ξ^c and ξ^i are independent standard Brownian motions, independent of the Brownian motions B^c and $\{B^j\}$ of (5).

The risk-neutral distribution of $(\lambda^{*,i}, \lambda^i)$ is specified by assuming that

$$\begin{aligned} \sqrt{\rho} dB_t^c + \sqrt{1 - \rho} dB_t^i &= -\frac{\kappa \theta^i - \tilde{\kappa} \tilde{\theta}^i}{\sigma} dt - \frac{\tilde{\kappa} - \kappa}{\sigma} X_t^i dt \\ &\quad + \sqrt{\rho} d\tilde{B}_t^c + \sqrt{1 - \rho} d\tilde{B}_t^i, \end{aligned} \quad (18)$$

and that

$$\begin{aligned} \sqrt{\rho_u} d\xi_t^c + \sqrt{1 - \rho_u} d\xi_t^i &= -\frac{\kappa_u \theta_u}{\sigma_u} dt - \frac{\tilde{\kappa}_u - \kappa_u}{\sigma_u} u_t^i dt \\ &\quad + \sqrt{\rho_u} d\tilde{\xi}_t^c + \sqrt{1 - \rho_u} d\tilde{\xi}_t^i, \end{aligned} \quad (19)$$

where $\tilde{B}^c, \tilde{B}^i, \tilde{\xi}^c$, and $\tilde{\xi}^i$ are independent standard Brownians motion under the risk-neutral measure Q , independent of $\{\tilde{B}^j\}_{j \neq i}$ and $\{\tilde{\xi}^j\}_{j \neq i}$, and where $\tilde{\theta}^i, \tilde{\kappa}$, and $\tilde{\kappa}_u$ are constants. In addition to the parameter vector Θ , the model for λ^* requires an

estimator of the parameter vector

$$\Theta^* = (\beta_0, \beta_1, \beta_2, \{\tilde{\theta}^i\}, \tilde{\kappa}, \theta_u, \kappa_u, \sigma_u, \rho_u, \tilde{\kappa}_u).$$

Since we have limited data with which to fit the parameters in Θ^* , we impose the over-identifying restriction that $\tilde{\kappa} = \kappa$ and $\tilde{\theta}^i = \theta^i - \gamma$, where γ is a constant, the same for all firms in the sector. This model is overly restrictive with respect to the potential for differences between actual and risk-neutral dynamics. As a result, our ability to disentangle the various contributions to default risk premia is limited.

C. Estimation Strategy and Results

For any given sector, we estimate the parameters (Θ, Θ^*) for the joint model of actual and risk-neutral intensity processes in a two-step procedure. First, we estimate the sector EDF-implied parameter vector Θ of the actual intensity model λ , following the procedure described at the end of Section IV. We also compute the time series of quarterly market discount factors δ by bootstrapping the U.S. dollar-denominated LIBOR-quality swap yield curve using 3-, 6-, 9- and 12-month LIBOR rates, and 2-, 3-, 4-, and 5-year swap yields that we obtain from Datastream. In a second step, treating our estimate of Θ as though error-free, we estimate, sector by sector, the parameter vector Θ^* governing the risk-neutral intensity process λ^* . For this second step, our data consist of weekly (Wednesday) observations of 1-year and 5-year default swap rates and 1-year EDFs, from June 2000 through December 2004. As with the actual default intensity model, this is not a routine MLE procedure since the evaluation of the likelihood function requires a numerical differentiation of the modeled CDS rate $C(t)$ determined by (12), which we approximate using (14). At this point, we only use transition between consecutive matched pairs of CDS-EDF observations for which the EDF is not censored at 20%. This introduces a potential selection bias that we

are currently investigating, but do not expect to be significant.

We further break the estimation of the parameter vector Θ^* into three parts, in order to simplify the estimation procedure and to obtain more robust parameter estimates. First, assuming that $u_t^i = 0$, we determine β_0 , β_1 , β_2 , and γ so that the sum of squared differences between the logarithm of the observed 1-year and 5-year CDS rates and the logarithm of their model-implied counterparts is minimized. Second, assuming that 5-year CDS rates are measured without error, we impose over-identifying restrictions, in the form of moment conditions, by restricting Θ^* so that the model-implied stationary mean of $\exp(u)$ is 1, and so that the model-implied stationary mean of u is equal to the model-implied sample mean across all firms in a given sector. This improves the interpretability of the parameter estimates, and facilitates comparison of the implied values for λ and λ^* . We also choose σ_u and $\tilde{\kappa}_u$ so that the implied sample mean and sample standard deviation of the standardized innovations $\epsilon_{t+h}^i, \epsilon_{t+2h}^i, \dots$ of u_t^i , defined by

$$u_{t+h}^i = \theta_u + e^{-\kappa_u h} (u_t^i - \theta_u) + \sigma_u \sqrt{\frac{1 - e^{-2\kappa_u h}}{2\kappa_u}} \epsilon_{t+h}^i. \quad (20)$$

are close (in terms of mean squared differences) to 0 and 1, respectively. Finally, we find ρ_u so that the sector log-likelihood, using all previously determined parameters, is maximized. In all cases, λ and λ^* are measured in basis points per year. We assume that the risk-neutral mean loss given default L^* is 75% on average, across the three industries used in our study, and that the sector-specific levels of L^* are proportionally adjusted by the sector-specific average recoveries reported in Footnote 8.

Sector-by-sector parameter estimates for the oil-and-gas, healthcare, and broadcasting-and-entertainment industry are summarized in Table III, and sector-by-sector sample moments of the estimated risk premia, that is, the ratio of estimated risk-neutral to

Table III: Sector CDS-implied risk-neutral default intensity parameter estimates

parameter estimates	Oil and Gas	Healthcare	Broadcasting and Entertainment
$\hat{\beta}_0$	0.829	0.576	-2.685
$\hat{\beta}_1$	0.537	0.522	0.400
$\hat{\beta}_2$	0.707	0.628	1.594
$\hat{\gamma}$	0.563	-0.175	-0.421
$\hat{\theta}_u$	-0.645	-0.258	-0.283
$\hat{\kappa}_u$	0.290	0.197	0.248
$\hat{\sigma}_u$	0.864	0.451	0.530
$\hat{\rho}_u$	0.335	0.212	0.479
$\hat{\kappa}_u$	0.131	-0.200	-0.233
sector likelihood	0.851	1.436	1.388
L^*	0.646	0.836	0.768
no. firms	33	16	13

estimated actual default intensities, are provided in Table IV. Summary statistics by firm are listed in Table XIV, Appendix B.

As an illustrative example, Figure 8 displays the estimated ratio of risk-neutral to actual default probability for Disney, for each of several maturities. For the “instantaneous” maturity, this is the estimated jump-to-default risk premia, that is, the ratio of λ^* to λ . (The one-year risk-neutral and actual default probabilities are themselves individually plotted in Figure 1.) Figure 9 shows the observed 5-year CDS rates of Disney, as well as the contemporaneous CDS rates that would have been estimated by our model in the absence of any risk premia. Figure 9 also shows the CDS rates that would have applied in the absence of risk premia associated with non-default mark-to-market risk, that is, assuming no market price of risk associated with random fluctuations in the risk-neutral intensity processes, and taking all risk premia to be those associated with jump to default risk. (This is equivalent to artificially taking the risk-neutral distribution λ^* to be equal to the estimated empirical distribution of λ^* .)

Jarrow, Lando, and Yu (2005) provide conditions under which there are no jump-to-default risk premia ($\lambda = \lambda^*$). A sufficient condition, for example, is that there

Table IV: Sample moments for estimated jump-to-default risk premia, $\lambda_t^{*,i}/\lambda_t^i$.

	mean	median	min	max	1 st quartile	3 rd quartile
Oil and Gas	3.303	2.539	0.204	21.128	1.263	4.490
Healthcare	2.168	1.853	0.603	8.975	1.346	2.640
Broadcast.-E.	2.037	1.497	0.124	12.794	0.723	2.787
All	2.757	2.032	0.124	21.128	1.121	3.554

are infinitely many firms exposed to the same risk factors as the firm in question, all defaulting independently conditional on those risk factors. If $\lambda^* = \lambda$, then default risk premia are entirely due to the market price of risk for uncertainty in the adjustment of λ^* over time. Collin-Dufresne, Goldstein, and Hugonnier (2004) provide a theory for jump-to-default risk premia, based on a form of contagion. Collin-Dufresne, Goldstein, and Helwege (2004) provide some related empirical evidence. Even without contagion, jump-to-default risk premia can be large if it is difficult to hedge the risk associated with the timing of default, and loss given default, of a particular firm (except of course by directly transferring those risks to another investor). Given our assumption that the 1-year CDS rate is measured with noise, and given the relatively short time horizon that we use to estimate the market prices of risk associated with the Brownian motions driving risk-neutral intensities, we do not claim accuracy for our estimated decomposition of CDS risk premia into the portion due to jump-to-default risk (differences between λ^* and λ) and that due to the market prices of risk of factors driving changes over time in λ^* .

Consistent with the presence of market prices of risk associated with fluctuations in λ^* , these multiplicative risk premia shown in Figure 8 are generally larger for longer maturities. This term effect, associated with aversion to mark-to-market risk associated with changes in credit spreads, apparently dominates, empirically, a countervailing “convexity effect.”¹⁵

For example, extracting from Table III the fit implied for the healthcare sector,

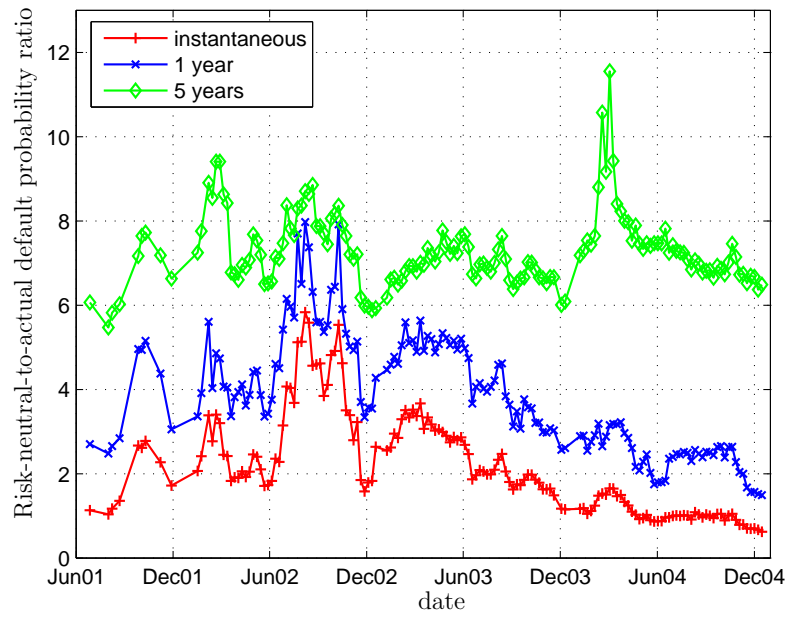


Figure 8: Estimated ratio of risk-neutral to actual default probabilities for Disney, by maturity.

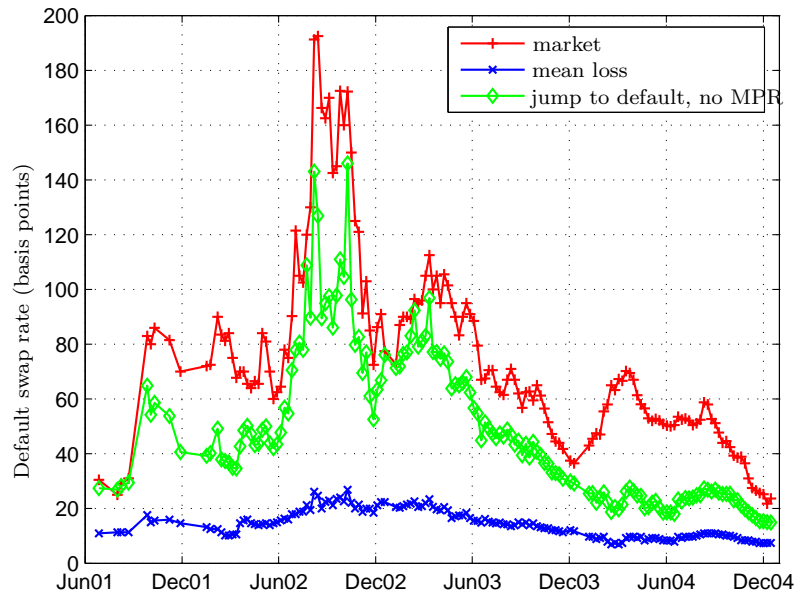


Figure 9: Disney: 5-year actual (market) CDS rates, and modeled CDS rates in the absence of any risk premia (mean loss), and in the absence of non-default mark-to-market risk premia.

we have

$$\log \lambda_t^* = 0.576 + 0.522 \log \lambda_t + 0.628 \log v_t + u_t,$$

or equivalently,

$$\lambda_t^* = 1.779 \lambda_t^{0.522} v_t^{0.628} e^{u_t},$$

where λ_t and λ_t^* are measured in basis points per year. So, for an actual default intensity of 100 basis points, a geometric average of all default intensities in the sector of 100 basis points, and $u_t = 0$, we get a risk-neutral default intensity of roughly 355 basis points. If the actual default intensity λ_t^i of firm i increases by 1%, then, everything else being equal, the risk-neutral default intensity $\lambda^{*,i}$ is estimated to increase by roughly $\beta_1\%$. Similarly, if the default intensities for each firm in the sector increase by 1%, $\lambda^{*,i}$ increases by roughly $(\beta_1 + \beta_2)\%$. The estimated risk-neutral distributions of λ and u are implied by the estimated model,

$$\begin{aligned} d \log \lambda_t &= 0.421 ((\hat{\theta}^i + 0.175) - \log \lambda_t) dt + 1.231 d\tilde{B}_t, \\ d \log u_t &= -0.200 u_t dt + 0.451 d\tilde{\xi}_t, \end{aligned}$$

where $\hat{\theta}^i$ is reported in Table XI, Appendix B. The sample averages of the estimated jump-to-default risk premia (that is, λ^*/λ) are 3.30, 2.17, and 2.04 for the oil-and-gas, healthcare, and broadcasting-and-entertainment sector, respectively. Additional sector-by-sector sample statistics of the estimated jump-to-default risk premia are provided in Table IV.

As a diagnostic check, we examine the behavior of the standardized innovations $\epsilon_{t+h}, \epsilon_{t+2h}, \dots$ of u_t , defined in (20). Under the specified model, and under the actual probability measure P , these innovations are standard normals. Table V lists the

Table V: Sample moments for standardized innovations

	Mean	SD
Oil and Gas	-0.034	0.982
Healthcare	-0.046	0.938
Broadcasting and Entertainment	-0.034	0.989
All	-0.037	0.975

sample mean and the sample standard deviation (SD) of the fitted versions of these standardized innovations, for each of the three sectors. Figure 10 shows the associated histogram of fitted ϵ_t , merging across all firms, plotted against the standard normal density curve. The innovations are relatively symmetrically distributed and somewhat leptokurtic.

Bearing in mind that our CDS rate observations are likely to be rather “noisy” relative to what actual market transaction rates would have been, as explained in Section 2, we undertook as a further robustness check an analysis of the implications of assuming that even the 5-year CDS rate is measured with error. This requires a filter for the underlying state variable u_t^i that, along with the observed EDF, determines what the “true” CDS would be. For this purpose, we assumed that the measurement noise of the CDS rate is such that the implied measurement noise for u_t^i is normally distributed, and *iid*. With this, because u_t^i is modeled as an auto-regressive Gaussian process, Kalman filtering applies. Some details of the main model were modified in order to make estimation tractable in the presence of filtering.¹⁶ Even so, this is not a simple procedure. (Details can be provided upon request.) Sector-by-sector maximum-likelihood parameter estimates for this model specification are provided in Table XV in the appendix. For the healthcare and oil-and-gas sectors, the estimated standard deviations of the measurement noise are approximately 5.5% and 5.8%, respectively. (These translate into roughly similar proportional levels of measurement noise for the associated CDS rates.) We do not yet have reliable measurement-noise results for the broadcasting-and-entertainment sector. Observed CDS rates do seem

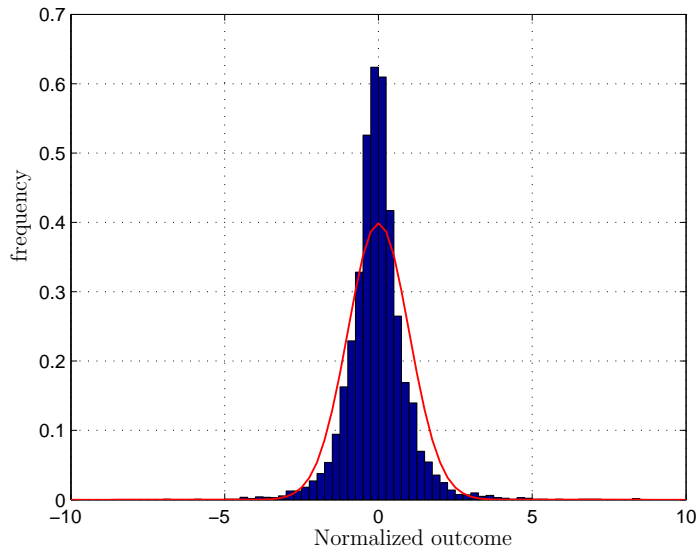


Figure 10: Estimated innovations ϵ across all sectors, and the standard normal density.

to best track sector-wide default risk in the broadcasting-and-entertainment sector.¹⁷ Measurement noise, when not treated in a simple autoregressive time-series model, causes an upward bias in estimated mean reversion coefficients. (For example, given an unusually large outcome of the current period’s measurement noise, the subsequent period’s measurement noise, being independent, has conditional mean zero, inducing an “extra” source of mean reversion in the observed time series.) Incorporating measurement noise into our model specification indeed caused large reductions in the estimated rates of mean reversion of u^i in the modified model, for both the healthcare and oil-and-gas sectors. In terms of the magnitudes and time variation of default risk premia, however, the broad characterizations that we draw on the basis of our main model are not dramatically affected.

VI. Discussion and Conclusion

We have uncovered substantial time variation in the compensation to investors for bearing U.S. corporate default risk. This section discusses some alternative explana-

tions for the time variation in default risk premia uncovered in our panel regression and time-series analyses. We will briefly explore three potential influences:

- Mismeasurement of actual conditional default probabilities.
- Time variation in risk-neutral conditional expectation of loss given default.
- Changes in the supply of and demand for risk bearing, whose effects are exaggerated by some limits on the mobility of capital across segments of the capital markets.
- The impact of principal-agency inefficiencies in the asset management industry.

We discuss these in order.

By construction, EDFs are unbiased estimates of default rates on average over their in-sample period. Suppose, however, that, like ratings, Moody's KMV EDFs are "too smooth" over time, so that they are biased downward when true conditional default probabilities are high, and biased upward when true conditional default probabilities are low. (In fact, Table II shows substantial annualized volatility, about 122%, for the default intensities implied by EDFs, belying the idea that EDFs vary little over time.) If EDFs were excessively smooth over time, then our estimated default risk premia would vary more dramatically than would actual default risk premia. Bohn, Arora, and Korablev (2005) report (in their Section 4.4) that EDFs did indeed predict "too many" defaults in 2003, a year in which we estimate declining default risk premia, as CDS rates came down faster than EDFs. For example, they estimate that the model that produces Moody's KMV EDFs would have placed a probability of only 27.3% on the event that there would have been as few or fewer defaults in 2003 by firms in their study sample than the actual number of defaults.¹⁸ If credit market participants had assigned lower default probabilities than the associated EDFs, then actual default risk premia for 2003 would be higher than those that we have estimated. On the other hand, the low number of realized defaults for 2003 could simply

have been a surprise (to anyone with accurate probability assignments). We are not aware, in any case, that marginal investors in corporate debt had access to better default probability estimates than those supplied by Moody’s KMV, but of course this is hard to verify. For the other years in our sample, the incidence of defaults was not especially “surprising” in this sense, relative to the EDF-predicted number of defaults. Bohn, Arora, and Korablev (2005) estimate the associated p -values for 2000, 2001, 2002, 2003, and 2004 at 46.1%, 61.8%, 47.9%, 27.3%, and 54.4%, respectively. In particular, the “ex-post” p -values for 2002 and 2004 are similar, but our estimated default risk premia are substantially higher in 2002 than in 2004. While our finding that default risk premia varied significantly during our sample period could be at least in part an artifact of mismeasured default probabilities, it is also not easy to make a strong case that EDFs were biased relative to the conditional default probabilities assigned by credit-market participants, time by time, in a manner that would largely explain our results.

A weakness of the methodology that we used to measure default risk premia is that it ignores correlation between the loss given default ℓ of an issuer and the default time τ . From Moody’s data covering default and default recoveries for all rated corporate debt from 1980 through 2004, a regression of cross-sectional average default recovery rate Y on the average default rate X provides the OLS estimated model $Y = 0.57 - 0.076X$, with an R^2 of 0.46, showing a highly significant and economically important negative relationship, in the aggregate, in recovery and default rates. Further analysis of this relationship is provided by Altman, Brady, Resti, and Sironi (2003). It might therefore be appropriate to assume that, for a given issuer, the LGD ℓ and the indicator of default before maturity, $1_{\tau < T}$, are positively correlated, risk-neutrally. If so, then our modeled CDS rates would be too low, for the risk-neutral mean default loss $E^Q(1_{\tau < T}\ell)$ is in that case larger than the product of $Q(\tau < T) = E^Q(1_{\tau < T})$ and the risk-neutral mean LGD, $L^* = E^Q(\ell)$. There are

currently no reliable data bearing on the magnitude of this effect. If this covariance effect is constant, one could scale the average effect of this covariance into the parameter assumed for the risk-neutral mean LGD. While the measured default risk premia could be biased from this effect, it does not necessarily follow that this correlation effect has a major impact on relative default risk premia at different times.

On the other hand, one might worry that the upward impact of this correlation on CDS rates is not constant over time, but greatest when default risk is highest, which could lead to an overstatement by our model (which ignores LGD-default correlation) of the time variation of default risk premia. For example, consider the extreme case in which there are no default risk premia (that is, $Q = P$). Suppose that, as we have found empirically, a default intensity λ_t is persistent over time. Let W_t denote the recovery that would occur in the event of default at time t . That is, $\ell = 1 - W_\tau$. Suppose that $W_t = f(\lambda_t, \epsilon)$, where ϵ is independent of the path of λ , and $f(x, y)$ is decreasing in x . Then, given persistence in λ_t , the conditional expectation $E_t(1 - W_\tau)$ of the LGD that will occur at the default time τ is increasing in the current intensity λ_t . The ratio of CDS rates to default probabilities would then be time-varying, and higher when CDS rates are higher. A model that ignores LGD-default correlation would misinterpret this as time variation in default risk premia.

Based on Moody's data for 1980 to 2004,¹⁹ Figure 11 shows the sample correlation between aggregate default rates in year t and average senior-unsecured debt recovery rates K years later, as the lag K ranges from 0 (contemporaneous) to 5 years, the maturity of our benchmark default swaps. Now, conditional on a default within the 5-year maturity, except for very low-quality firms, the expected time to default is roughly 2.5 years.²⁰ At least based on the data underlying this figure, there is no obvious reason to conclude that the LGD-default correlation effect on 5-year CDS rates is large. That said, there are almost no data bearing on the risk-neutral (as opposed to actual) LGD-default correlation, which would be needed to deduce the

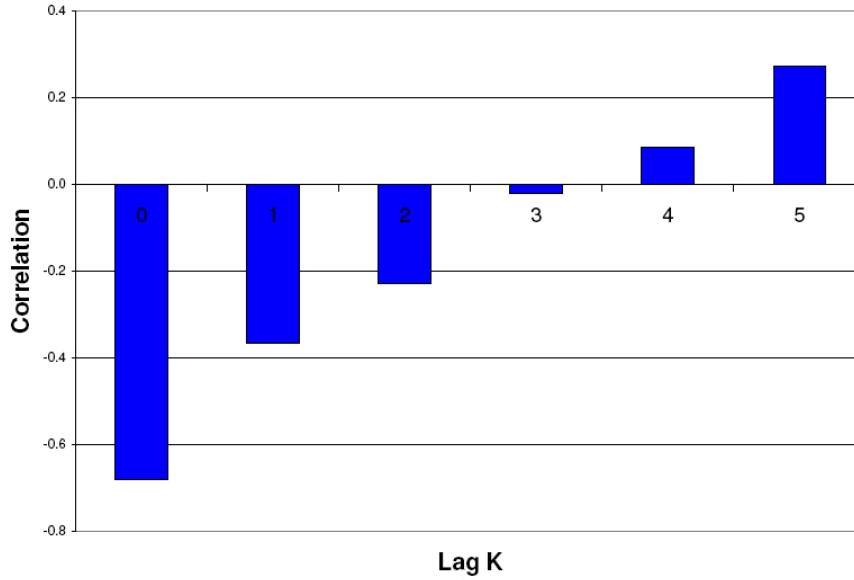


Figure 11: Correlation between default rate at year t and recovery rate at year $t+K$, for 1980-2004. Data from Moody's (2005).

impact of recovery risk on CDS rates.

In order to gauge the general magnitude of the effect of risk-neutral LGD-default correlation on CDS rates, we calculate the pricing impact in a simple example. Suppose, given the current risk-neutral intensity λ_t^* , that any default recovery W_t that occurs at time t is beta distributed²¹ with mean $M(\lambda_t^*)$ and standard deviation $0.2 - 0.4|0.5 - M(\lambda_t^*)|$. We assume a log-normal risk-neutral default intensity (as in our model), but for simplicity assume a 50% risk-neutral mean-reversion rate and 100% volatility. We set $M(\lambda_t^*) = e^{a-b\lambda_t^*}$, where a and b are chosen for an unconditional risk-neutral mean LGD of 0.5 and a risk-neutral correlation of -0.5 between $W_{2.5}$ and $\lambda_{2.5}^*$. (Again, 2.5 years is roughly the time horizon that matters for 5-year CDS, if one is to pick a particular time horizon.) At these parameters, an increase in default intensity from 200 basis points to 1000 basis points reduces current risk-neutral mean recovery from 50% to 25%. This is roughly the speculative-grade empirical experience from 1998 to 2001, a period during which average default recovery sank to an extremely low level, relative to history. Even before considering idiosyncratic recovery risk at the firm level that is partially washed out by aggregation, this represents a

substantial amount of LGD-default correlation, relative to our sample, most of which is of investment-grade quality. Moving the risk-neutral intensity-LGD correlation from 0 (the assumption in our CDS model) to 0.5 increases 5-year CDS rates from 100 basis points to about 118 basis points. This means that the risk-neutral default probability that would be inferred by a model assuming no LGD-default correlation, at an observed CDS rate of 118 basis points would be biased about 18% too high. This bias would always be positive, but would (under our distributional assumptions) be proportionately greater when CDS rates are high than when they are low. Such a bias is likely to be responsible for some of the measured variation in default risk premia reported in this paper. To repeat, whether the magnitude of the bias is large is difficult to judge because there are essentially no empirical data bearing directly on risk-neutral LGD-default correlation.

A third explanation for the large time variation in estimated default risk premia that we have uncovered in this paper may be based on the usual market-equilibrium suspect: variations over time in the supply of, and demand for, risk bearing, potentially exacerbated by limited mobility of capital across different classes of asset markets. Along the lines of the explanation suggested by Froot and O'Connell (1999) for time variation in catastrophe insurance risk premia, capital moves into and out of the market for corporate credit in response to fluctuations in risk premia, but perhaps not instantaneously so. Generally, when there are large losses or large increases in risk in a particular market segment, if capital does not move immediately out of other asset markets and into that segment, then risk premia would tend to adjust so as to match the demand for capital with the supply of capital that is available to the sector. Investors or asset managers with available capital take time to be found by intermediaries, to be convinced (perhaps being unfamiliar with the particular asset class) of the available risk premia, and to exit from the markets in which they are currently invested. For the catastrophe risk insurance market, Froot and O'Connell

(1999) show that this process can take well over a year, in terms of the half-life of the mean reversion of risk premia to long-run levels. Similar explanations, albeit with shorter half-lives, have been offered by Gabaix, Krishnamurthy, and Vigneron (2007) for variation of prepayment risk premia in the market for mortgage-backed securities, and by Greenwood (2005) for the price impact of supply shocks in equity markets.

In order to explore the role of limited capital mobility in determining credit risk premia, we replaced the time and sector dummies in the panel-regression (2) with current stock-market volatility V (in percent), as measured by VIX (an index of option-implied volatility of the S-and-P 500), and with the total face value D of U.S. defaulted corporate debt over the prior 6 months, measured in billions of U.S. dollars. (The defaulted debt data were provided by Moody's.) As market volatility goes up, a given level of capital available to bear risk represents less and less capital per unit of risk to be borne. If replacement capital does not move into the corporate debt sector immediately, the supply and demand for risk capital will match at a higher price per unit of risk. (This effect would be present even with perfect capital mobility, but the magnitude of the effect is increased with partially segmented markets.) Similarly, a loss of capital through trailing defaulted debt, proxied by D , reduces the amount of capital available to bear risk. The fitted model and “robust” standard errors (shown parenthetically) are

$$\log \text{CDS}_i = 1.08 + 0.84 \log \text{EDF}_i + 1.18D + 0.011V + z_i, \quad (21)$$

$$(0.015) \quad (0.003) \quad (0.57) \quad (0.00049)$$

where z_i is the residual. The associated R^2 is 0.71. The coefficients for trailing defaulted debt and VIX are statistically significant at conventional confidence levels, particularly so for VIX, whose coefficient has a t -statistic of over 22. The EDF-based default probability estimates already incorporate the impact of volatility on

default risk through an estimate of the contemporaneous volatility of each firm, so the estimated effect of a change in VIX on CDS rates is on top of that implied by the impact of a change in market volatility on estimated default probabilities.²² (It is of course possible that volatility is not well estimated in the Moody's KMV EDF model.) The estimated coefficient for VIX implies that the reduction in S-and-P 500 volatility that occurred between August 2002 and the end of 2003, from 45% to about 11%, is associated with a proportional reduction of about 44% in the credit default swap rates assigned by the market at a given default probability. This is plausible, in that the time fixed effects associated with (2) are, on average across the three sectors, of similar magnitude. The role of trailing defaulted debt, while statistically significant, is somewhat less pronounced in magnitude. In July 2002, for example, trailing-6-month corporate debt increased over that of the previous month by 3.4 billion dollars, which is responsible for an estimated proportional increase in credit spreads of approximately 4%, holding all else equal. On the other hand, if the EDF model is imperfect, a change in trailing defaulted debt could proxy for a change in average default probabilities not captured by the EDF model. We cannot rule this out. In light of the rather adverse market conditions of mid 2002, a behavioral reaction by some market participants also cannot be ruled out.

On top of these effects, money managers may have been reluctant to place themselves in "harm's way," in terms of adverse inference by investors regarding the ability or efforts of asset managers in light of prior losses through default. This principal-agency effect may have reduced their willingness to load up on corporate credit risk, despite the high risk premia offered. A related principal-agency explanation of declining risk premia during 2002-2004 is the propensity for fixed-income asset managers to "reach for yield" when treasury market rates decline, as they did during 2002-2004. That is, in order to offer their supposedly unsophisticated or poorly informed investors "fixed-income" yields that do not decline markedly as treasury rates decline,

money managers are willing to take increasing credit risk for the same yield. While reaching for yield is frequently mentioned anecdotally, we have no specific evidence of its prevalence.

Yet another interpretation of the estimated role of market-wide volatility and trailing defaulted debt is that these variables are proxies for an increase in default correlation, for which we have not controlled. Variation over time in conditional default correlation can be responsible for increases or reductions in the degree of diversification available in the corporate bond market, and therefore could change default risk premia. Capital immobility would magnify any such effect.

Consistent with our conjecture that variation in default risk premia are partly caused by sluggish movement in risk capital across sectors, Collin-Dufresne, Goldstein, and Martin (2001) earlier showed that VIX is an important explanatory variable for changes in credit spreads, after controlling, firm by firm, with equity returns. They did not pin down an explanation for the role of VIX. For different data, the same important role of VIX was confirmed by Schafer and Strebulaev (2008). Notably, Collin-Dufresne, Goldstein, and Martin (2001) emphasize, and Avramov, Jostova, and Philipov (2004), Schafer and Strebulaev (2008), and Yu (2006) confirm, that a large fraction of the variation in a firm's credit spreads is *not* explained by the same firm's equity returns. From the theoretical view that debt and equity can be treated as derivatives written on the total market value of the underlying firm, perfect capital mobility between equity and debt markets would tend to lead debt and equity returns to have strong common components, contrary to the results of these studies. Consistent with our conjecture that there may be risk premia in the corporate debt sector that are due to temporary fluctuations in the availability of risk capital due to partially segmented markets, Collin-Dufresne, Goldstein, and Martin (2001) find statistical evidence of a further common factor in corporate bond returns, above and beyond equity returns, risk-free yields, and VIX, whose source was unex-

Table VI: Five-year default risk premium implied by the structural-model results of Huang and Huang (2003)

Initial Rating	Premium (ratio)	$Q(\tau < 5)$ (percent)	$P(\tau < 5)$ (percent)
Aaa	1.7497	0.04	0.02
Aa	1.7947	0.09	0.05
A	1.7322	0.25	0.15
Baa	1.4418	1.22	0.84
Ba	1.1658	9.11	7.85
B	1.1058	25.61	23.41

plained. Although Avramov, Jostova, and Philipov (2004) do not find support for this “mysterious common factor,” Saita (2005) does.

The market did indeed respond over time to opportunities for insuring default risk. Open interest in the default swap market has roughly doubled in each year for the last several years. Investment banks and broker dealers in corporate credit markets have increased their credit-derivatives staffs and market-making capacity significantly. Among hedge funds that were newly established during the period 2002-2004, those specializing in corporate credit risk were by far the most prevalent.

Our results on the average magnitude of default risk premia are comparable to those available in the prior literature. Using the structural model of Leland and Toft (1996), Huang and Huang (2003) calibrated parameters for the model determining actual and risk-neutral default probabilities, by credit rating, that are implied from equity-market risk premia, recoveries, initial leverage ratios, and average default frequencies. All underlying parameters were obtained from averages reported by the credit rating agencies, Moody’s and Standard and Poors, except for the equity-market risk premia, which were obtained by rating from estimates by Bhandari (1999). At the five-year maturity point, our calculation of the associated estimated ratios of annualized risk-neutral to actual five-year default probabilities are reported in Table VI. In magnitude, the results are also roughly consistent with those of Driessen (2005). In both of these prior studies, risk premia are proportionately higher for highly-rated

firms, consistent with the results of our panel-regression model and time-series analyses.

Appendices

A. MLE for Intensity from EDFs

This appendix shows our methodology for MLE estimation of the parameters of the default intensity, including the effects of missing EDF data as well as censoring of EDFs by truncation from above at 20%. Our data is the monthly observed EDF level Y_i at month i , for each of $N + 1$ month-end times t_0, t_1, \dots, t_N .

From (4), for any time t and time step h (which is 1/12 in our application), the discretely sampled log-intensity process X satisfies

$$X_{t+h} = b_0 + b_1 X_t + \epsilon_{t+h}, \quad (\text{A.1})$$

where $b_1 = e^{-\kappa h}$, $b_0 = \theta(1 - b_1)$, and $\epsilon_{t+h}, \epsilon_{t+2h}, \dots$ are *iid* normal with mean zero and variance $\sigma_\epsilon = \sigma^2(1 - e^{-2\kappa h})/(2\kappa)$.

For a given firm, we initialize the search for the parameter vector $\Theta = (\theta, \kappa, \sigma)$ as follows. First, we regress $\log(Y_i)$ on $\log(Y_{i-1})$, using only months at which both the current and the lagged EDF are observed and not truncated at 20%. The associated regression coefficient estimates, denoted by \hat{b}_0 and \hat{b}_1 , are considered to be starting estimates of b_0 and b_1 , respectively. The sample standard deviation of the fitted residuals, $\hat{\sigma}_\epsilon$, is our starting estimate for σ_ϵ . We then start the search for $\Theta = (\theta, \kappa, \sigma)$ at

$$\begin{aligned} \kappa_0 &= -\frac{\log(\hat{b}_1)}{h}, \\ \theta_0 &= \frac{\hat{b}_0}{1 - \hat{b}_1}, \\ \sigma_0 &= \hat{\sigma}_\epsilon \sqrt{\frac{2\kappa_0}{1 - \exp(-2\kappa_0 h)}}. \end{aligned}$$

If Θ is the true parameter vector, then $Y_i = G(\lambda(t_i); \Theta)$, where G is defined via (5).

Suppose, to pick an example of a censoring outcome from which the general case can easily be deduced, that for months $k + 1$ through $\bar{k} \geq k + 1$, inclusive, the EDFs are truncated at $\zeta = 20\%$, meaning that the censored and observed EDF is 20%, implying that the actual EDF was larger than or equal to 20%, and moreover that the EDF data from month $l + 1 > \bar{k} + 1$ to month \bar{l} are missing. Let $\mathcal{I} = \{i : k + 1 \leq i \leq \bar{k}\} \cup \{i : l + 1 \leq i \leq \bar{l}\}$ denote the censored and missing month numbers. Then the likelihood of the observed non-censored EDFs $Y = \{Y_i : i \notin \mathcal{I}\}$ evaluated at outcomes $y = \{y_i : i \notin \mathcal{I}\}$, using the usual abuse of notation for measures, is defined by

$$\begin{aligned} \mathcal{L}(Y, \mathcal{I}; \Theta) dy &= \prod_{n=0}^{k-1} P(Y_{n+1} \in dy_{n+1}; Y_n = y_n, \Theta) \\ &\quad \times P(Y_{k+1} \geq \zeta, \dots, Y_{\bar{k}} \geq \zeta, Y_{\bar{k}+1} \in dy_{\bar{k}+1}; Y_k = y_k, \Theta) \\ &\quad \times \prod_{n=k+1}^{l-1} P(Y_{n+1} \in dy_{n+1}; Y_n = y_n, \Theta) \\ &\quad \times P(Y_{l+1} \in dy_{l+1}; Y_l = y_l, \Theta) \\ &\quad \times \prod_{n=l+1}^{N-1} P(Y_{n+1} \in dy_{n+1}; Y_n = y_n, \Theta), \end{aligned}$$

where $P(\cdot; Y_n = y_n; \Theta)$ denotes the distribution of $\{Y_{n+1}, Y_{n+2}, \dots\}$ associated with initial condition y_n for Y_n , and associated with parameter vector Θ . A maximum likelihood estimator (MLE) $\hat{\Theta}$ for Θ solves

$$\sup_{\Theta} \mathcal{L}(Y, \mathcal{I}; \Theta). \tag{A.2}$$

For $z \in \mathbb{R}$, we let $g(z; \Theta) = G(e^z; \Theta)$, and let $Z_i^\Theta = g^{-1}(Y_i; \Theta)$ denote the logarithm of the default intensity at time t_i that would be implied by a non-censored EDF observation Y_i , assuming the true parameter vector is Θ . Letting $Dg(\cdot; \Theta)$ denote

the partial derivative of $g(\cdot; \Theta)$ with respect to its first argument, and using standard change-of-measure arguments, we can rewrite the likelihood as

$$\begin{aligned}
 \mathcal{L}(Y, \mathcal{I}; \Theta) &= \prod_{n=0}^{k-1} P(Z_{n+1}^\Theta; Z_n^\Theta, \Theta) (Dg(Z_{n+1}^\Theta; \Theta))^{-1} \\
 &\quad \times P(Y_{k+1} \geq \zeta, \dots, Y_{\bar{k}} \geq \zeta; Y_k = y_k, Y_{\bar{k}+1} = y_{\bar{k}+1}, \Theta) \\
 &\quad \times P(Z_{\bar{k}+1}^\Theta; Z_{\bar{k}}^\Theta, \Theta) (Dg(Z_{\bar{k}+1}^\Theta; \Theta))^{-1} \\
 &\quad \times \prod_{n=\bar{k}+1}^{l-1} P(Z_{n+1}^\Theta; Z_n^\Theta, \Theta) (Dg(Z_{n+1}^\Theta; \Theta))^{-1} \\
 &\quad \times P(Z_{l+1}^\Theta; Z_l^\Theta, \Theta) (Dg(Z_{l+1}^\Theta; \Theta))^{-1} \\
 &\quad \times \prod_{n=\bar{l}+1}^{N-1} P(Z_{n+1}^\Theta; Z_n^\Theta, \Theta) (Dg(Z_{n+1}^\Theta; \Theta))^{-1}. \tag{A.3}
 \end{aligned}$$

The second term on the right-hand side of (A.3) is equal to

$$\begin{aligned}
 q(Y; \Theta) &= P(Z_{k+1}^\Theta \geq g^{-1}(\zeta; \Theta), \dots, Z_{\bar{k}}^\Theta \geq g^{-1}(\zeta; \Theta); \\
 &\quad Z_k^\Theta = g^{-1}(y_k; \Theta), Z_{\bar{k}+1}^\Theta = g^{-1}(y_{\bar{k}+1}; \Theta), \Theta).
 \end{aligned}$$

In the remainder of this appendix, we describe how to compute $q(Y; \Theta)$ by Monte Carlo integration, and hence $P(Y_{k+1} \geq \zeta, \dots, Y_{\bar{k}} \geq \zeta; Y_k = y_k, Y_{\bar{k}+1} = y_{\bar{k}+1}, \Theta)$. In order to simplify notation we suppress Θ in what follows. We observe that for any time t between times s and u , the conditional distribution of $X(t)$ given $X(s)$ and $X(u)$ is a normal distribution with mean $M(t|s, u)$ and variance $V(t|s, u)$ given by

$$\begin{aligned}
 M(t|s, u) &= \frac{1 - e^{-2\kappa(u-t)}}{1 - e^{-2\kappa(u-s)}} M(t|s) + \frac{e^{-2\kappa(u-t)} - e^{-2\kappa(u-s)}}{1 - e^{-2\kappa(u-s)}} M(t|u), \\
 V(t|s, u) &= \frac{V(t|s)V(u|t)}{V(u|s)},
 \end{aligned}$$

where, for times t before u , we let

$$\begin{aligned} M(u | t) &= \theta + e^{-\kappa(u-t)}(X(t) - \theta) \\ V(u | t) &= \frac{\sigma^2}{2\kappa}(1 - e^{-2\kappa(u-t)}) \\ M(t | u) &= e^{\kappa(u-t)}(X(u) - \theta(1 - e^{-\kappa(u-t)})) \end{aligned}$$

denote the conditional expectation and variance, respectively, of $X(u)$ given $X(t)$, and the conditional expectation of $X(t)$ given $X(u)$. As a consequence, letting $Z_k = X(t_k)$, we can easily simulate from the joint conditional distribution of $(Z_{k+1}, \dots, Z_{\bar{k}})$ given Z_k and $Z_{\bar{k}+1}$, which is given by

$$\begin{aligned} P(Z_{k+1}, \dots, Z_{\bar{k}} | Z_k, Z_{\bar{k}+1}) &= P(Z_{k+1} | Z_k, Z_{\bar{k}+1}) \\ &\quad \times \prod_{j=1}^{\bar{k}-(k+1)} P(Z_{k+j+1} | Z_{k+j}, Z_{\bar{k}+1}). \end{aligned}$$

We are now in a position to estimate the quantity in (A.4) by generating some “large” integer number J of independent sample paths $\{(Z_{k+1}^j, \dots, Z_{\bar{k}}^j); 1 \leq j \leq J\}$ from the joint conditional distribution of $(Z_{k+1}, \dots, Z_{\bar{k}})$ given Z_k and $Z_{\bar{k}+1}$, and by computing the fraction of those paths for which $Z_i^j \geq g^{-1}(\zeta)$ for all i in $\{k+1, \dots, \bar{k}\}$.

Table VII: Number of observations of 1-year EDFs. Data: Moody's KMV.

Ticker	sector [†]	not		capped at		total	Ticker	sector	not		capped at		total
		censored	20%	0.02%	20%				censored	0.02%	20%		
253647Q	B&E	55	18	0	73	IPG	B&E	119	1	0	120		
ABC	H	113	0	0	113	JNJ	H	26	118	0	144		
ABT	H	82	62	0	144	KMG	O&G	123	21	0	144		
ADELQ	B&E	97	0	16	113	KMI	O&G	142	2	0	144		
AGN	H	126	18	0	144	KMP	O&G	139	3	0	142		
AHC	O&G	137	7	0	144	KRI	B&E	70	50	0	120		
AMGN	H	36	108	0	144	L	B&E	81	21	0	102		
APA	O&G	144	0	0	144	LH	H	120	0	0	120		
APC	O&G	144	0	0	144	LLY	H	101	43	0	144		
BAX	H	144	0	0	144	MCCC	B&E	56	0	0	56		
BC	B&E	120	0	0	120	MDT	H	39	105	0	144		
BEV	H	142	0	2	144	MGLH	H	103	0	19	122		
BHI	O&G	144	0	0	144	MMM	H	42	78	0	120		
BJS	O&G	144	0	0	144	MRK	H	50	70	0	120		
BLC	B&E	115	5	0	120	MRO	O&G	144	0	0	144		
BMY	H	54	90	0	144	NBR	O&G	144	0	0	144		
BR	O&G	129	15	0	144	NEV	O&G	137	0	0	137		
BSX	H	123	21	0	144	NOI	O&G	97	0	0	97		
CAH	H	144	0	0	144	OCR	H	131	13	0	144		
CAM	O&G	113	0	0	113	OEI	O&G	124	0	0	124		
CCU	B&E	142	2	0	144	OMC	B&E	120	0	0	120		
CHIR	H	144	0	0	144	OXY	O&G	132	12	0	144		
CHK	O&G	130	0	13	143	PDE	O&G	144	0	0	144		
CHTR	B&E	50	0	11	61	PFE	H	36	84	0	120		
CMCSA	B&E	144	0	0	144	PHA	H	77	23	0	100		
CNG	U	72	13	0	85	PKD	O&G	144	0	0	144		
COC	O&G	45	0	0	45	PRM	B&E	107	0	3	110		
COP	O&G	133	11	0	144	PXD	O&G	144	0	0	144		
COX	B&E	116	0	0	116	RCL	B&E	141	0	0	141		
CVX	O&G	39	105	0	144	RIG	O&G	136	4	0	140		
CYH	H	97	0	0	97	SBGI	B&E	113	0	0	113		
DCX	A	68	6	0	74	SGP	H	61	59	0	120		
DGX	H	93	0	0	93	SLB	O&G	85	35	0	120		
DIS	B&E	103	41	0	144	SUN	O&G	120	0	0	120		
DO	O&G	97	10	0	107	THC	H	144	0	0	144		
DVN	O&G	135	9	0	144	TLM	O&G	126	18	0	144		
DYN	U	121	0	13	134	TRI	H	67	0	0	67		
EEP	O&G	114	6	0	120	TSG	B&E	94	3	0	97		
ENRNQ	O&G	105	1	1	107	TSO	O&G	135	0	0	135		
EP	O&G	143	0	1	144	TWX	B&E	144	0	0	144		
EPD	O&G	76	0	0	76	UCL	O&G	117	3	0	120		
F	A	143	1	0	144	UHS	H	120	0	0	120		
FST	O&G	142	0	0	142	UNH	H	113	7	0	120		
GDT	H	115	2	0	117	VIA	B&E	139	5	0	144		
GENZ	H	144	0	0	144	VLO	O&G	144	0	0	144		
GLM	O&G	120	0	0	120	VPI	O&G	144	0	0	144		
GM	A	144	0	0	144	WFT	O&G	143	1	0	144		
HAL	O&G	144	0	0	144	WLP	H	137	7	0	144		
HCA	H	140	4	0	144	WMB	U	135	1	8	144		
HCR	H	120	0	0	120	WYE	H	86	58	0	144		
HMA	H	95	25	0	120	XOM	O&G	0	120	0	120		
HRC	H	131	0	4	135	XTO	O&G	120	0	0	120		
HUM	H	142	0	0	142	YBTVA	B&E	120	0	0	120		
ICCI	T	65	0	0	65								

[†] A: Automobile; B&E: Broadcasting and Entertainment; H: Healthcare; O&G: Oil and Gas; R: Retail; T: Transportation; U: Utilities.

B. Additional Background Statistics

This appendix contains additional background statistics regarding the firms studied. Section II contains the data regarding firms from the broadcasting-and-entertainment industry. This appendix includes information regarding the firms studied from the healthcare and the oil-and-gas industries.

Table VIII: Healthcare firms

Firm Name	Median EDF	Median Rating	No. Quotes
Abbott Laboratories	4.0	A1	1,845
Allergan Inc	3.0	A3	2,137
Amerisource Bergen Corp	83.5	Ba3	437
Amgen Inc	2.0	A2	2,159
Baxter International Inc	32.0	Baa1	2,252
Beverly Enterprises Inc	1,086.0	B1	285
Boston Scientific Corp	5.0	Baa1	1,813
Bristol-Myers Squibb Co	22.0	A1	2,063
Cardinal Health Inc	15.0	Baa3	1,753
Chiron Corp	12.0	Baa2	1,920
Community Health Systems Inc	98.0	N/A	307
Eli Lilly & Co	3.0	Aa3	1,942
Genzyme Corp	24.0	N/A	1,657
Guidant Corp	5.0	Baa1	1,407
HCA Inc	23.0	Ba2	891
Health Management Associates Inc	10.0	N/A	2,222
Healthisouth Corp	–	N/A	318
Humana Inc	40.0	Baa3	1,925
Johnson & Johnson	2.0	Aaa	1,654
Laboratory Corp Of America Holdings	12.0	Baa3	1,635
Manor Care Inc	21.0	Baa3	1,168
Medtronic Inc	2.0	N/A	2,093
Merck & Co Inc	5.0	Aa2	1,516
Minnesota Mining & Manufacturing Co (3M)	2.0	Aa1	1,655
Pfizer Inc	2.0	Aaa	1,504
Pharmacia Corporation	9.0	Aaa	1,116
Quest Diagnostics	10.0	Baa2	1,230
Schering-Plough Corporation	25.0	Baa1	1,658
Tenet Healthcare Corporation	67.0	B3	–
Triad Hospitals Inc	148.0	B2	519
United Health Group Inc	2.0	A3	1,442
Universal Health Services Inc	33.5	Baa3	1,237
Wellpoint Health Networks	–	Baa1	1,580
Wyeth	17.0	Baa1	2,150

Table IX: Oil and gas firms

Firm Name	Median EDF	Median Rating	No. Quotes
Amerada Hess Corp	20.0	Ba1	1,284
Anadarko Petroleum Corp	43.0	Baa1	2,696
Apache Corp	11.0	A3	2,217
Baker Hughes Inc	15.0	A2	2,207
BJ Services Co	17.0	Baa2	1,588
BurlingtonResourcesInc	10.0	Baa1	2,056
Chesapeake Energy Corp	177.0	Ba3	1,152
Chevron Texaco Corp	3.0	N/A	1,897
Conoco Phillips Holding Co	15.5	A3	1,677
Cooper Cameron Corp	29.0	Baa1	1,518
XTO Energy Inc	6.0	Baa3	1,225
Diamond Offshore Drilling	25.0	Baa2	2,331
EL Paso Corp	1,000.0	Caa1	2,294
Exxon Mobil Corp	2.0	N/A	1,329
Forest Oil Corp	107.0	Ba3	367
Global Marine Inc	12.0	Baa1	1,401
Halliburton Co	86.0	Baa2	2,139
Kerr-Mc Gee Corp	38.0	Baa3	2,170
Kinder Morgan Energy Partners LP	12.0	Baa1	2,263
Kinder Morgan Inc	7.0	Baa2	2,003
National-Oilwell Inc	31.0	Baa2	1,201
Occidental Petroleum Corp	8.0	Baa1	2,581
Parker Drilling Co	446.5	B2	449
Conoco Phillips	6.0	A3	2,929
Pioneer Natural Resources Co	44.0	Baa3	1,001
Pride International Inc	113.5	Ba2	1,228
Shell Oil Co	-	Aa2	1,373
Sunoco Inc	7.5	Baa2	1,536
Talisman Energy Inc	5.0	N/A	1,425
Transocean Inc	71.5	Baa2	2,487
Unocal Corp	6.0	Baa2	1,441
Marathon Oil Corp	10.0	Baa1	2,024
Valero Energy Corp	35.5	Baa3	2,637
Vintage Petroleum Inc	229.0	Ba3	556
Weatherford International Ltd	35.0	Baa1	2,874
Enron Corp	51.5	WR	361
Devon Energy Corporation	19.0	Baa2	2,878
Enterprise Products Partners LP	5.0	N/A	1,439
Enbridge Energy Partners LP	6.0	N/A	1,192
Nabors Industries Ltd	41.0	N/A	2,594
Schlumberger Ltd	8.0	N/A	1,673
Schlumberger Technology Corp	-	A2	1,021
Duke Energy Field Services Llc	-	Baa2	1,004

Table X: Estimated log-CDS panel regression (2), using daily median CDS data.

	Estimate [†]	Std. Error	Estimate [†]	Std. Error	Estimate [†]	Std. Error
Number of CDS Samples	33,912					
Intercept	1.448	0.047				
Slope	0.760	0.015				
Sector-month dummy	Oil and Gas		Broadcasting and E.		Healthcare	
Dec-00	0.124	0.006	0.256	0.007	0.204	0.011
Jan-01	0.425	0.011	0.562	0.009	0.360	0.008
Feb-01	0.605	0.004	0.481	0.010	0.111	0.007
Mar-01	0.621	0.005	0.458	0.010	0.326	0.009
Apr-01	0.389	0.009	0.495	0.013	0.478	0.009
May-01	0.260	0.012	0.502	0.019	0.618	0.009
Jun-01	0.280	0.014	0.223	0.014	0.075	0.008
Jul-01	-0.044	0.015	0.145	0.016	0.019	0.010
Aug-01	-0.075	0.021	0.085	0.014	0.144	0.005
Sep-01	0.174	0.032	0.254	0.020	0.582	0.005
Oct-01	0.256	0.019	0.316	0.019	0.641	0.008
Nov-01	0.416	0.015	0.406	0.014	0.128	0.008
Dec-01 and Jan-02	0.227	0.017	0.352	0.020	0.465	0.011
Feb-02	0.306	0.020	0.551	0.023	0.508	0.021
Mar-02	0.229	0.018	0.582	0.021	0.478	0.021
Apr-02	0.209	0.016	0.576	0.024	0.424	0.015
May-02	0.270	0.015	0.661	0.024	0.542	0.014
Jun-02	0.200	0.016	0.651	0.024	0.396	0.009
Jul-02	0.162	0.018	0.970	0.030	0.487	0.012
Aug-02	0.278	0.017	1.005	0.031	0.556	0.012
Sep-02	0.171	0.018	0.861	0.029	0.332	0.011
Oct-02	0.253	0.019	0.830	0.031	0.427	0.012
Nov-02	0.175	0.018	0.712	0.029	0.412	0.013
Dec-02	0.157	0.016	0.522	0.028	0.418	0.014
Jan-03	0.125	0.016	0.348	0.027	0.320	0.015
Feb-03	0.126	0.015	0.269	0.026	0.217	0.014
Mar-03	0.060	0.014	0.251	0.024	0.015	0.010
Apr-03	-0.048	0.014	0.232	0.021	-0.004	0.009
May-03	-0.019	0.014	0.202	0.020	0.009	0.006
Jun-03	0.004	0.012	0.219	0.018	0.062	0.006
Jul-03	-0.143	0.012	0.131	0.015	0.035	0.007
Aug-03	-0.132	0.010	0.095	0.013	-0.036	0.005
Sep-03	-0.188	0.010	-0.068	0.013	-0.070	0.004
Oct-03	-0.251	0.011	-0.109	0.013	-0.073	0.004
Nov-03	-0.340	0.010	-0.193	0.009	-0.179	0.004
Dec-03	-0.245	0.007	-0.293	0.007	-0.228	0.003
Jan-04	-0.115	0.006	0.041	0.006	-0.135	0.001
Feb-04	0.052	0.004	0.241	0.006	-0.044	0.000
Mar-04	0.089	0.003	0.193	0.006	0.023	0.001
Apr-04	0.053	0.003	0.144	0.005	0.022	0.001
May-04	0.106	0.002	0.144	0.005	0.116	0.001
Jun-04	0.134	0.002	0.100	0.006	0.106	0.001
Jul-04	0.156	0.002	0.102	0.006	0.056	0.002
Aug-04	0.130	0.001	0.224	0.006	0.131	0.002
Sep-04	0.123	0.001	0.176	0.006	0.072	0.002
Oct-04	0.069	0.002	0.141	0.005	0.066	0.003
Nov-04	0.045	0.001	-0.013	0.004	0.130	0.002
Dec-04	reference		-0.175	0.004	-0.001	0.002
Sum of Squared Residuals	8,742.760					
Total Sum of Squares	34,098.286					
R^2	0.744					

[†] Regressions are based on CDS data for the period December 2000 through December 2004.

Table XI: Fitted parameters of actual default intensity models

Ticker	$\hat{\theta}$	$\hat{\kappa}$	$\hat{\sigma}$	Ticker	$\hat{\theta}$	$\hat{\kappa}$	$\hat{\sigma}$
253647Q	†	–	–	IPG	4.547	0.169	1.007
ABC	4.061	1.980	2.521	JNJ	†	–	–
ABT	†	–	–	KMG	2.231	0.239	0.910
ADELQ	5.869	0.539	1.541	KMI	1.593	1.670	2.753
AGN	1.284	0.366	0.902	KMP	2.537	0.383	1.041
AHC	2.215	0.574	1.116	KRI	†	–	–
AMGN	†	–	–	L	3.225	0.269	1.086
APA	2.211	0.298	0.909	LH	2.216	0.169	1.307
APC	2.033	0.238	0.855	LLY	†	–	–
BAX	2.360	0.649	1.148	MCCC	6.124	1.034	1.522
BC	2.761	0.377	1.082	MDT	†	–	–
BEV	4.740	0.586	1.424	MGLH	6.636	0.125	1.059
BHI	1.850	0.202	0.793	MMM	†	–	–
BJS	2.897	0.730	1.310	MRK	†	–	–
BLC	2.316	0.248	1.073	MRO	2.232	0.359	0.897
BMY	†	–	–	NBR	2.936	1.080	1.559
BR	1.897	0.401	0.994	NEV	4.562	0.274	0.955
BSX	1.813	0.701	1.822	NOI	3.571	1.196	1.900
CAH	2.183	0.595	1.293	OCR	2.445	0.245	1.311
CAM	3.167	0.592	1.101	OEI	4.192	0.331	1.227
CCU	2.445	0.390	1.483	OMC	2.755	1.209	1.769
CHIR	2.575	0.617	1.222	OXY	-0.378	0.078	0.696
CHK	3.261	0.167	1.265	PDE	4.234	0.811	1.533
CHTR	11.718	0.116	1.062	PFE	†	–	–
CMCSA	3.317	0.528	0.972	PHA	†	–	–
CNG	†	–	–	PKD	5.741	0.141	1.201
COC	2.702	1.959	1.965	PRM	6.693	0.054	1.142
COP	†	–	–	PXD	3.777	0.437	1.311
COX	2.345	0.647	1.387	RCL	2.661	0.382	1.210
CVX	†	–	–	RIG	2.216	0.313	1.377
CYH	4.285	0.957	1.478	SBGI	5.183	0.677	1.500
DCX	3.502	0.680	1.353	SGP	2.861	0.149	0.571
DGX	0.408	0.154	0.856	SLB	1.846	0.289	0.871
DIS	1.773	0.360	0.879	SUN	2.452	0.307	0.933
DO	1.882	0.298	1.437	THC	3.526	0.393	1.002
DVN	2.274	0.335	1.392	TLM	1.886	0.278	1.200
DYN		–	–	TRI	4.189	0.656	0.842
EEP	1.657	0.194	0.867	TSG	3.678	0.241	1.255
ENRNQ	‡	–	–	TSO	4.097	0.544	1.299
EP	5.014	0.264	1.040	TWX	3.253	0.296	1.097
EPD	1.731	1.377	2.666	UCL	1.506	0.188	0.833
F	2.568	0.401	1.127	UHS	3.039	0.880	1.168
FST	4.478	0.825	1.345	UNH	1.786	0.302	1.235
GDT	1.618	0.562	1.083	VIA	2.133	0.657	1.452
GENZ	2.309	0.817	1.486	VLO	2.682	0.281	1.038
GLM	2.283	0.307	1.025	VPI	4.047	0.751	1.547
GM	3.008	0.974	1.358	WFT	2.132	0.189	1.102
HAL	2.967	0.407	1.457	WLP	2.646	0.740	1.469
HCA	2.004	0.427	1.740	WMB	3.699	0.211	1.258
HCR	2.679	0.361	1.016	WYE	1.812	0.536	0.875
HMA	2.115	0.259	0.961	XOM	†	†	†
HRC	4.033	0.398	1.668	XTO	2.299	0.155	0.978
HUM	3.936	0.370	1.390	YBTVA	5.440	0.793	1.535
ICCI	6.164	0.610	1.402				

† No estimates provided; the sample mean of the 1-year EDF is less than 10 basis points.

‡ No estimates within admissible parameter region; the estimate for the mean-reversion parameter κ is negative.

|| Firm removed from data set.

Table XII: MC distribution of default intensity parameter estimates

	$\hat{\theta}$	$\hat{\kappa}$	$\hat{\sigma}$
“true parameters”	4.00	0.50	1.00
10 years of data			
mean	3.95	0.87	1.21
std. dev.	(0.50)	(0.40)	(0.26)
50 years of data			
mean	3.98	0.58	1.04
std. dev.	(0.21)	(0.15)	(0.09)

Table XIII: Fitted parameters of default intensity models

Oil and Gas		Healthcare		Broadcasting and E.	
Ticker	$\hat{\theta}$	Ticker	$\hat{\theta}$	Ticker	$\hat{\theta}$
AHC	2.092	ABC	4.542	ADELQ	5.862
APA	2.863	AGN	1.775	BC	3.063
APC	2.410	BAX	2.289	BLC	2.248
BHI	2.674	BEV	4.701	CCU	2.485
BJS	3.230	BSX	1.974	CHTR	6.359
BR	1.707	CAH	2.320	CMCSA	3.436
CAM	3.429	CHIR	2.765	COX	2.479
CHK	5.379	DGX	3.926	DIS	1.503
COC	2.861	GDT	2.189	IPG	3.338
DO	2.632	GENZ	2.546	L	2.175
DVN	2.587	HCA	2.778	MCCC	5.943
EEP	2.094	HCR	2.805	OMC	2.883
ENRNQ	1.915	HMA	2.059	PRM	4.417
EP	3.647	HRC	4.025	RCL	3.219
EPD	2.529	HUM	3.667	SBGI	5.021
FST	4.863	LH	3.916	TSG	2.915
GLM	3.173	MGLH	5.789	TWX	3.278
HAL	3.260	OCR	2.959	VIA	2.289
KMG	2.198	SGP	1.400	YBTVA	5.324
KMI	2.328	THC	3.459		
KMP	2.459	TRI	4.709		
MRO	2.727	UHS	2.930		
NBR	3.168	UNH	2.162		
NEV	4.449	WLP	2.570		
NOI	3.942	WYE	1.447		
OEI	3.889				
OXY	2.489				
PDE	4.513				
PKD	4.363				
PXD	3.890				
RIG	2.896				
SLB	1.804				
SUN	2.743				
TLM	2.191				
TSO	4.376				
UCL	2.102				
VLO	3.021				
VPI	4.137				
WFT	3.630				
XTO	3.767				

Table XIV: Summary statistics by firm for risk-neutral default intensity models

Oil and Gas				Healthcare				Broadcasting and Entertainment			
Ticker	in J ?	mean(λ^*/λ)	median(λ^*/λ)	Ticker	in J ?	mean(λ^*/λ)	median(λ^*/λ)	Ticker	in J ?	mean(λ^*/λ)	median(λ^*/λ)
AHC	1	4.978	5.042	AGN	1	3.885	3.800	BC	0	1.084	1.096
APA	1	1.809	1.760	BAX	1	1.254	0.879	BLC	0	2.561	2.479
APC	1	1.516	1.362	BSX	1	2.755	2.741	CCU	1	2.164	1.635
BHI	1	1.250	0.963	CAH	1	1.823	1.817	CMCSA	1	3.139	1.491
BJS	1	2.028	1.006	CHIR	1	1.784	1.751	COX	0	3.673	3.327
BR	1	3.933	3.848	DGX	0	1.695	1.827	DIS	1	2.097	1.852
CAM	1	0.595	0.544	GDT	0	2.697	2.518	IPG	0	0.915	0.865
CHK	1	2.637	2.733	GENZ	0	3.397	2.092	L	0	2.112	1.495
COC	0	4.769	4.791	HCR	0	1.674	1.606	OMC	0	1.481	1.459
DO	1	1.927	1.763	HMA	0	2.515	2.523	RCL	0	2.128	1.819
DVN	1	4.119	3.263	HUM	1	1.083	1.045	TSG	0	0.496	0.296
EEP	0	6.696	6.899	LH	0	1.780	1.781	TWX	1	0.914	0.610
EP	0	2.673	1.837	SGP	0	1.240	1.078	VIA	1	2.351	1.924
EPD	0	14.622	14.599	UHS	0	1.272	1.252				
GLM	0	1.184	1.104	UNH	0	4.222	3.944				
HAL	1	3.375	2.185	WYE	1	1.915	1.812				
KMG	1	2.790	2.590								
KMI	1	8.009	7.280								
KMP	0	3.564	3.342								
MRO	0	2.396	2.431								
NBR	1	1.162	0.896								
NOI	0	0.535	0.451								
OXY	0	5.169	5.108								
PDE	1	4.712	4.450								
PXD	1	7.650	8.328								
RIG	0	1.188	1.104								
SLB	0	1.659	1.560								
SUN	0	3.257	3.072								
TLM	1	5.707	5.407								
UCL	0	2.741	2.701								
VLO	1	4.761	3.883								
WFT	1	0.828	0.765								
XTO	0	7.523	7.542								

Table XV: Sector CDS-implied Kalman-filter-based risk-neutral default intensity parameter estimates for the model specification with measurement noise for both 1-year and 5-year CDS

parameter estimates [†]	Oil and Gas	Healthcare	Broadcasting and Entertainment
$\hat{\beta}_0$	3.083	1.296	1.469
$\hat{\beta}_1$	0.164	0.424	0.133
$\hat{\beta}_2$	0.077	0.169	0.546
$\hat{\kappa}_u$	0.423	1.052	0.761
$\hat{\sigma}_u$	1.198	1.991	1.860
$\hat{\rho}_u$	0.394	0.248	0.688
mean $\hat{\theta}^i$	3.769	1.451	2.456
$\hat{\kappa} = \hat{\kappa}_u$	0.350	0.377	0.350
SD (measurement noise)	0.058	0.055	0.000
sector likelihood	1.090	0.797	0.825
L^*	0.646	0.836	0.768
no. firms	33	16	13

[†] The parameters are estimated using MLE. We assume that $\tilde{\kappa} = \tilde{\kappa}_u$, and restrict γ and θ_u to be zero. For each firm i , we determine $\tilde{\theta}^i$ so that the model-implied average 1-year CDS rate equals the observed average rate.

REFERENCES

- Altman, E., B. Brady, A. Resti, and A. Sironi (2003). The Link between Default and Recovery Rates: Theory, Empirical Evidence and Implications, forthcoming, *Journal of Business*. Working Paper, Stern School of Business New York University NY.
- Artzner, P. and F. Delbaen (1992). Credit Risk and Prepayment Option. *ASTIN Bulletin* **22**, 81–96.
- Avramov, D., G. Jostova, and A. Philipov (2004). Credit Risk Changes: Common Factors and Firm-Level Fundamentals. Working Paper, University of Maryland.
- Behar, R. and K. Nagpal (1999). Dynamics of Rating Transition. Working Paper, Standard and Poor's.
- Bhandari, C. (1999). Debt/Equity Ratio and Expected Common Stock Returns: Empirical Evidence. *Journal of Finance* **43**, 507–528.
- Bharath, S. and T. Shumway. Forecasting Default with the Merton Distance-to-Default Model. *Review of Financial Studies*.
- Black, F. and P. Karasinski (1991). Bond and Option Pricing when Short Rates are Log-Normal. *Financial Analysts Journal*, 52–59.
- Black, F. and M. Scholes (1973). The Pricing of Options and Corporate Liabilities. *Journal of Political Economy* **81**, 637–654.
- Blanco, R., S. Brennan, and I. Marsh (2005). An Empirical Analysis of the Dynamic Relationship between Investment Grade Bonds and Default Swaps. *Journal of Finance* **60**, 2255–2282.
- Bohn, J. (2000). An Empirical Assessment of a Simple Contingent-Claims Model for the Valuation of Risky Debt. *Journal of Risk Finance*, Summer, 55–77.
- Bohn, J., N. Arora, and I. Korablev (2005). Power and Level Validation of the EDF Credit Measure in the U.S. Market. Working Paper, Moody's KMV, San Francisco.
- Cao, C., F. Yu, and K. Zhong (2007). The Information Content of Option-Implied Volatility for Credit Default Swap Valuation. Working Paper, Penn State University, State College.
- Chen, L., P. Collin-Dufresne, and R. Goldstein (2005). On the Relationship between Credit Spread Puzzles and the Equity Premium Puzzle. Working Paper, University of California, Berkeley.
- Collin-Dufresne, P., R. Goldstein, and J. Helwege (2004). Is Credit Event Risk Priced? Modeling Contagion via the Updating of Beliefs. Working Paper, University of California, Berkeley.
- Collin-Dufresne, P., R. S. Goldstein, and J. Hugonnier (2004). A General Formula for Valuing Defaultable Securities. *Econometrica* **72**, 1377–1408.

- Collin-Dufresne, P., R. S. Goldstein, and J. S. Martin (2001). The Determinants of Credit Spread Changes. *Journal of Finance* **56**, 2177-207.
- Crosbie, P. J. and J. R. Bohn (2002). Modeling Default Risk. KMV, LLC.
- Delbaen, F. and W. Schachermayer (1999). A General Version of the Fundamental Theorem of Asset Pricing. *Mathematische Annalen* **300**, 463–520.
- Delianedis, G. and R. Geske (1998). Credit Risk and Risk Neutral Default Probabilities: Information About Rating Migrations and Defaults. Working Paper, Working Paper 19-98, Anderson Graduate School of Business, University of California, Los Angeles.
- Delianedis, G., R. Geske, and T. Corzo (1998). Credit Risk Analysis with Option Models: Estimation and Comparison of Actual and Risk-Neutral Default Probabilities. Working Paper, Working Paper, Anderson Graduate School of Business, University of California, Los Angeles.
- Driessen, J. (2005). Is Default Event Risk Priced in Corporate Bonds? *Review of Financial Studies* **18**, 165–195.
- Duffee, G. (1998). The Relation Between Treasury Yields and Corporate Bond Yield Spreads. *Journal of Finance* **53**, 2225–2242.
- Duffie, D. (1999). Credit Swap Valuation. *Financial Analysts Journal*, January–February, 73–87.
- Duffie, D. (2001). *Dynamic Asset Pricing Theory (third edition)*. Princeton University Press.
- Duffie, D., L. Saita, and K. Wang (2007). Multiperiod Corporate Default Probabilities with Stochastic Covaraites. *Journal of Financial Economics* **83**, 635–665.
- Duffie, D. and K. Singleton (2003). *Credit Risk*. Princeton University Press.
- Eom, Y., J. Helwege, and J.-Z. Huang (2004). Structural Models of Corporate Bond Pricing: An Empirical Analysis. *Review of Financial Studies* **17**, 499–544.
- Fischer, E., R. Heinkel, and J. Zechner (1989). Dynamic Capital Structure Choice: Theory and Tests. *Journal of Finance* **44**, 19–40.
- Fisher, L. (1959). Determinants of the Risk Premium on Corporate Bonds. *Journal of Political Economy* **67**, 217–237.
- Fons, J. (1987). The Default Premium and Corporate Bond Experience. *Journal of Finance* **42**, 81–97.
- Froot, K. and P. O’Connell (1999). The Pricing of U.S. Catastrophe Reinsurance. In K. Froot (Ed.), *The Financing of Catastrophe Risk*, pp. 195–232. University of Chicago Press.
- Gabaix, X., A. Krishnamurthy, and O. Vigneron (2007). Limits of Arbitrage: Theory and Evidence from the Mortgage-Backed Securities Market. *Journal of Finance* **62**, 557-95.

- Greenwood, R. (2005). Short- and Long-Term Demand Curves for U.S. Stocks: Theory and Evidence on the Dynamics of Arbitrage. *Journal of Financial Economics* **75**, 607–649.
- Harrison, M. and D. Kreps (1979). Martingales and Arbitrage in Multiperiod Securities Markets. *Journal of Economic Theory* **20**, 381–408.
- Huang, J.-Z. and M. Huang (2003). How Much of the Corporate-Treasury Yield Spread is Due to Credit Risk: A New Calibration Approach. Working Paper, Penn State University.
- Hull, J. and A. White (1994). Numerical Procedures for Implementing Term Structure Models I: Single Factor Models. *Journal of Derivatives* **2**, 7-16.
- Jarrow, R. A., D. Lando, and F. Yu (2005). Default Risk and Diversification: Theory and Applications. *Mathematical Finance* **15**, 1–26.
- Kavvathas, D. (2001). Estimating Credit Rating Transition Probabilities for Corporate Bonds. Working Paper, University of Chicago.
- Kealhofer, S. (2003). Quantifying Credit Risk I: Default Prediction. *Financial Analysts Journal*, January–February, 30–44.
- Kurbat, M. and I. Korbalev (2002). Methodology for Testing the Level of the EDF Credit Measure. Working Paper, Technical Report 020729, Moody’s KMV.
- Kusuoka, S. (1999). A Remark on Default Risk Models. *Advances in Mathematical Economics* **1**, 69–82.
- Lando, D. (1998). Cox Processes and Credit-Risky Securities. *Review of Derivatives Research* **2**, 99–120.
- Lando, D. and T. Skødeberg (2002). Analyzing Rating Transitions and Rating Drift with Continuous Observations. *Journal of Banking and Finance* **26**, 987–1019.
- Leland, H. and K. Toft (1996). Optimal Capital Structure, Endogenous Bankruptcy, and the Term Structure of Credit Spreads. *Journal of Finance* **51**, 987–1019.
- Longstaff, F., S. Mithal, and E. Neis (2005). Corporate Yield Spreads: Default Risk or Liquidity? New Evidence from the Default Swap Market. *Journal of Finance* **60**, 213–2254.
- Merton, R. (1974). On The Pricing of Corporate Debt: The Risk Structure of Interest Rates. *The Journal of Finance* **29**, 449–470.
- Nickell, P., W. Perraudin, and S. Varotto (2000). Stability of Ratings Transitions. *Journal of Banking and Finance* **24**, 203–228.
- Saita, L. (2005). The Puzzling Price of Corporate Credit Risk (Preliminary). Working Paper, Stanford University.
- Schafer, S. and I. Strebulaev (2008). Structural Models of Credit Risk are Useful: Evidence from Hedge Ratios on Corporate Bonds. forthcoming, *Journal of Financial Economics*.

- Vassalou, M. and Y. Xing (2004). Default Risk in Equity Returns. *Journal of Finance* **59**, 831–868.
- Wooldridge, J. (2002). *Econometric Analysis of Cross Section and Panel Data*. Cambridge, Mass.: MIT Press.
- Yu, F. (2006). How Profitable is Capital Structure Arbitrage? *Financial Analysts Journal* **62**, 47–62.

Notes

¹The one-year accuracy ratio, a traditional measure of accuracy for default prediction, is $2 \int_0^1 [f(x) - x] dx$, where $f(x)$ is the fraction of the firms that defaulted within one year after the time of prediction that are within the lowest scoring fraction x of firms, according to the predictive method.

²In a 2004 report, “High Yield Credit Default Swaps,” Fitch Ratings reports substantially higher recoveries for 2001-2003 for restructurings (52.7%) than for default by missed payment (29.4%) or bankruptcy (25.3%).

³We used the following algorithm: (a) If a bid and an ask were present, we record the bid-ask spread. (b) If the bid is missing, we subtract the average bid-ask spread to estimate the ask. (c) If the ask is missing, we add the average bid-ask spread to estimate the bid. (d) From the resulting bid and ask, we calculate the mid-quote as the average of the bid and ask quotes.

⁴Under iid measurement noise, more precise estimates would be obtained by using all CDS observations separately, rather than the median CDS observation. We prefer using the median, given the potential damage caused by outliers.

⁵We also examined the fit, by non-linear least squares, of the model, $Y_i = \alpha X_i^\beta + u_i$, which differs from (1) by having a residual that is additive in levels, rather than additive in logs. An informal comparison shows that the non-linear least-squares model is somewhat preferred for lower-quality firms.

⁶For an EDF of 10 basis points, the model (2) implies a fitted CDS rate of $e^{1.45+0.76 \times \log 10+0.52^2/2}$, accounting for the effect of normality of disturbances, and using the fact that $E(e^X) = e^{\text{var}(X)/2}$ for a zero-mean normal random variable X .

⁷For example, for the oil and gas sector, the dummy coefficients for August 2002 and December 2003 are 0.278 and -0.245 , respectively, for a proportionate change in fitted CDS rates at a given 5-year EDF of $e^{-0.245-0.278} - 1 = -0.41$.

⁸ From the Moody’s sectoral data, the average recovery for the oil-and-gas sector

is estimated from the simple average of the of the Moody's "Oil and Oil Services" and the "Utility-Gas" sectors, at 48%. Broadcasting and Entertainment recoveries are estimated at the 'Media Broadcasting and Cable' average of 38%, and Healthcare at 32.7%.

⁹See Black and Karasinski (1991).

¹⁰We explored more tractable affine jump-diffusion specifications, but the fitted short-horizon conditional sample variances of changes in intensity varied in a manner much closer to linear in the square of intensity than to constant-plus-linear in the level of intensity, as would be dictated by affine models.

¹¹Details are available from the authors upon request. The Matlab code can be downloaded from the web site www.andrew.cmu.edu/user/aberndt/software/.

¹²Here, r is a progressively measurable process with $\int_0^t |r(s)| ds < \infty$ for all t , such that there exists a "money-market" trading strategy, allowing investment at any time t of one unit of account, with continual re-investment until any future time T , with a final value of $e^{\int_t^T r(s) ds}$.

¹³A more precise definition of L_t^* is given on page 130 of Duffie and Singleton (2003).

¹⁴This ignores the impact of a default event on the time-series properties of v , which is small provided the influence of any one firm on the geometric average is small. In our data, J includes those firms marked with "1" in Appendix C, which are essentially the largest and most liquidly traded firms. As it happens, none of these defaulted during our sample period. Details are provided in Table XIV in the appendix.

¹⁵The risk-neutral survival probability $E^Q \left(e^{-\int_0^T \lambda^*(t) dt} \right)$ is larger by Jensen's Inequality than $e^{-\int_0^T E^Q[\lambda^*(t)] dt}$. Suppose that $\lambda(t)$ is constant for simplicity, and consider the natural assumption that the unconditional variance of $\lambda^*(t)/\lambda(t)$ grows with t . Then, even if $E^Q[\lambda^*(t)/\lambda(t)]$ does not depend on t , the ratio of the risk-neutral to the actual probability of default by t would typically decline with maturity. This effect, however, is apparently more than offset, empirically, for example by trends in

$E(\lambda^*(t))$ or by market prices of risk associated with random changes in λ^* .

¹⁶We replace the restriction that $\tilde{\kappa} = \kappa$ by the assumption that $\tilde{\kappa} = \tilde{\kappa}_u$. Moreover, we lift the restriction of a constant market-price-of-risk parameter γ , and instead determine $\tilde{\theta}^i$ for each firm i so that the model-implied average 1-year CDS rate is equal to the observed average rate.

¹⁷For this model specification, $\hat{\beta}_2 = 0.546$ for the broadcasting-and-entertainment sector, compared to $\hat{\beta}_2 = 0.169$ and $\hat{\beta}_2 = 0.077$ for the healthcare and the oil-and-gas sectors, respectively.

¹⁸For this study, Bohn, Arora, and Korablev (2005) considered U.S. public firms with assets of at least \$300 million dollars. Out of the total sample of 1594 firms at the beginning of 2003, 12 defaulted.

¹⁹These data are available at moodys.com.

²⁰The conditional mean default time converges to 2.5 years as the default probability converges to zero, at a constant default intensity.

²¹Moody's KMV uses the beta distribution for its modeled recovery distributions in its LossCalc model.

²²Cao, Yu, and Zhong (2007) explore the ability of firm-level option-implied volatility to explain CDS rate changes.