Measuring Default Risk Premia from Default Swap Rates and EDFs*

Antje Berndt,[†] Rohan Douglas,[‡] Darrell Duffie,[§]
Mark Ferguson,[¶] and David Schranz[‖]

First Version: April 16, 2003.
Current Version: March 20, 2004

Abstract

This paper estimates recent default risk premia for U.S. corporate debt, based on a close 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 22 banks and specialty dealers, allow us to establish a strong link between actual and risk-neutral default probabilities for the 69 firms 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 thrid quarter of 2002, dropping by roughly 50% to late 2003.

*We 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 are also grateful to Gustavo Manso and Leandro Saita for research assistance, and to Linda Bethel and Sandra Berg for technical assistance.

[†]Department of Operations Research and Industrial Engineering, Cornell University.
[‡]Quantifi LLC, New York, New York.
[§]Graduate School of Business, Stanford University.
[¶]Quantifi LLC, New York, New York.
[‖]CIBC, Toronto.

1 Introduction

This paper estimates recent default risk premia for U.S. corporate debt, based on a close relationship between default probabilities, as estimated by the Moody's KMV EDF measure, and default swap (CDS) market rates. The default-swap data, obtained by CIBC from 22 banks and specialty dealers, allow us to establish a strong link between actual and risk-neutral default probabilities for the 69 firms in the three sectors that we analyzed: broadcasting and entertainment, healthcare, and oil and gas.

Based on over 49,000 CDS rate quotes, we find that 5-year EDFs explain over 70% of the cross-sectional variation in 5-year CDS rates, after controlling for sectoral and for temporal effects. We find that the marginal impact of default probability on credit spreads is significantly greater for high-credit-quality firms than for low-credit-quality firms. We also find that, for a given default probability, there is 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 o n average roughly 50% lower than at their peak, after controlling for sectoral effects.

If a firm's risk-neutral default intensity λ^* and risk-neutral expected fraction L^* of notional lost at default are assumed to be relatively stable over time, the firm's CDS rate and its par-coupon credit spread would be approximately equal to the risk-neutral mean loss rate, $\lambda^* L^*$, ignoring illiquidity effects.[1]

The lowest 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 L^*, our measured relationship between CDS and EDF implies that risk-neutral default intensities are roughly double actual default intensities (proxied by EDFs), on average, although this premium is much higher for high quality

[1] This close relationship between risk-neutral mean loss rate and par credit spreads is from Duffie and Singleton (1999). The close relationship between par credit spreads and CDS rates is explained by Duffie (1999). With illiquidty, however, Duffie and Singleton (1999) show a divergence between spreads and risk-neutral mean loss rates. Longstaff, Mithal, and Neis (2003) claim that there are indeed illiquidity effects, measuring yield spreads relative to Treasury yields. Blanco, Brennan, and Marsh (2003) suggest that the illiquidity differences causing bond spreads and default swap rates to diverge are small if using interest-rate swap yields as a benchmark for bond spreads.

firms, and lower for low quality firms. This ratio of risk-neutral to actual default intensities is a default-timing risk premium whose size and behavior over time is a primary objective of our analysis.

This simple estimate of the default risk premium does not consider: *(i)* the effect of random fluctuations in actual and risk-neutral default intensities, *(ii)* the potential impact of illiquidity on CDS rates, *(iii)* variation between actual and risk-neutral mean fractional losses given default, *(iv)* correlation between fluctuations over time in risk-neutral mean losses given default and risk-neutral default intensities, *(v)* the effect of cheapest-to-deliver settlement options on default swap rates, and *(vi)* sample noise. We shall address the impact of each of these later in the paper.

Fons (1987) gave the earliest empirical analysis, to our knowledge, of the relationship between actual and risk-neutral default probabilities. By using *ex-post* default rates and excess returns on corporate debt, he was able to show that, for his data set treating the early 1980s, that market implied risk-neutral default rates were about 5% larger than actual default rates.

Driessen (2002) recently estimated the relationship between actual and risk-neutral default probabilities, using U.S. corporate bond price data (rather than CDS data), and using average long-horizon default frequencies by credit rating (rather than contemporaneous firm-by-firm EDFs). Driessen reported an average risk premium across his data of 1.89, after accounting for tax and liquidity effects, that is roughly in line with the estimates that we provide here. While the conceptual foundations of Driessen's study are similar to ours, there are substantial differences in our respective 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. Most importantly, we do not rely on historical average default rate by credit rating as a proxy for current conditional default intensity.

Because the corporate bonds in Driessen's study involve taxable coupon income, extracting credit spreads required an estimation by him of the portion of the bond yield spread that is associated with taxes. As for the estimated actual default probabilities, Driessen's reliance on average frequency of default for bonds of the same rating rules out conditioning on current market

conditions, which Kavvathas (2001) and others have shown to be significant. Reliance on default frequency by rating also rules out consideration of distinctions in default risk among bonds of the same rating. 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)). Finally, while Longstaff, Mithal, and Neis (2003) and Blanco, Brennan, and Marsh (2003) show that bond yield spreads and CDS rates provide roughly contemporaneous information, our enquiries of market participants have led us to the view that default swaps, because they are "un-funded exposures," in the language of dealers, have rates that are less sensitive to liquidity effects than are bond yield spreads.

Fisher (1959) took a simple regression approach to explaining yield spreads on corporate debt in terms of various credit-quality and liquidity related variables.

Bohn (2000), Delianedis and Geske (1998), G. Delianedis Geske and Corzo (1998), and Huang and Huang (2000) 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.

The potential applications of our study are numerous, and include: *(i)* the relationship between risk and expected return for the credit component of corporate debt, and *(ii)* analysis of the extent to which the default-risk premia of different firms have common factors, as well as the dynamics and macroeconomics of these common factors. These applications can, in turn, be further applied to a range of pricing and portfolio investment decisions involving corporate credit risk.

An example of our results is illustrated in Figure 1, which shows estimated actual and risk-neutral 1-year default probabilities for Vintage Petroleum, based on EDF and CDS rate data from Moody's KMV and CIBC, respectively. Figure 1 shows the typical pattern in our sample of high default risk premium in the third quarter of 2002.

The remainder of the paper is structured as follows. Section 2 describes our data, including a brief introduction to default swaps and to the construction of the Moody's KMV EDF measure of default probability. Section 3 presents simple descriptive statistical evidence of a strong relationship between CDS rates and EDFs across several sectors. Section 4 introduces a simple approximate time-series model of actual default intensities, and a maximum-likelihood approach to parameter estimation. Section 4 also con-

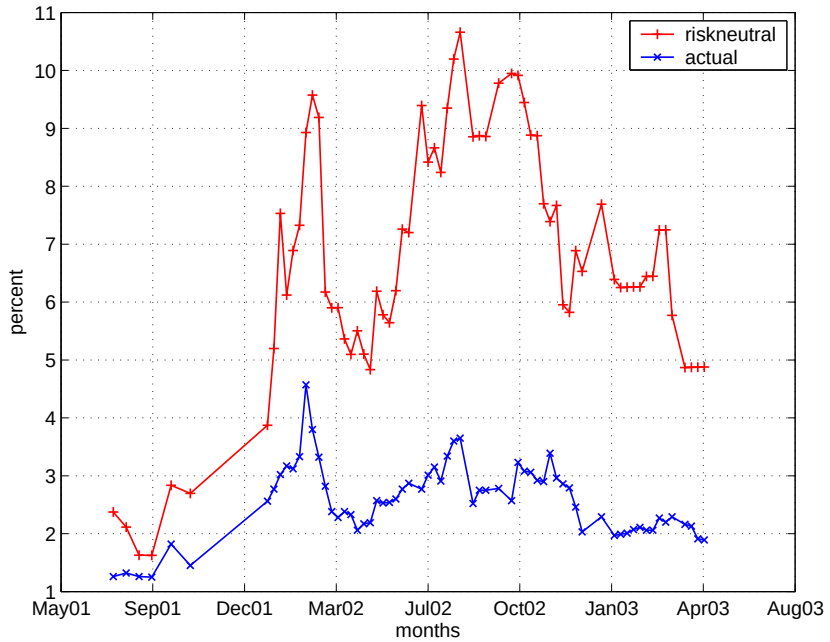

Figure 1: Estimated 1-year default probabilities for Vintage Petroleum. Sources: Moody's KMV and CIBC.

tains parameter estimates for each firm, based on 10 years of monthly observations of 1-year EDFs for each firm. Section 5 provides a reduced-form pricing model for default swaps, based on time-series models of actual and risk-neutral default intensities. Section 5.2 introduces our parameterization of the time-series model for risk-neutral default intensities, using both EDFs and CDS rates. Section 5.3 provides estimates of the parameters for each of the three sectors. Section 6 discusses the results, and then concludes.

2 The EDF and CDS Data

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

2.1 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 fitted 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, roughly speaking, the number of standard deviations of annual asset growth by which its current assets exceed a measure of book liabilities. The liability measure is, in the current implementation of the EDF model, equal to 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 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. While one could criticize the EDF measure as an estimator of the "true" conditional default probability, it has a number of important merits for business practice and for our study, relative to other available approaches to estimating conditional default probabilities. First, it 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.) Second, while the EDF model is based on a single covariate, distance-to-default, for default prediction, and one might wish to exploit additional covariates (Duffie and Wang (2003), Shumway (2001)), the distance-to-default (DD) covariate has a strong underlying theoretical basis in the Black-Scholes-Merton model, within which DD is a sufficient statistic for conditional default probabilities.

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

An alternative 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 Credit Metrics 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 (2000), Kavvathas (2001), Nickell, Perraudin, and Varotto (2000), among others.

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, it is the only widely used name-specific major source of conditional default probability 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. In order to develop a time-series model of default intensities, however, we turn in Section 4 to 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 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.

2.2 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, less the default-swap premium that has accrued since the last default-swap payment date. 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.

The CDS rate is the annualized premium rate, as a fraction of notional. 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. The inclusion, or not, of restructuring as a covered default event has been a question of debate among the community of buyers and sellers of protection. ISDA, the industry coordinator of standardized OTC contracts (www.isda.org), has arranged a consensus for a standardized contractual definition of default that, we believe, is likely to be reflected in most of our data. This consensus definition of default has been adjusted over time, 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 im-

plied default probabilities, since of course a wider definition of default implies a higher risk-neutral default probability.

If restructuring is included as a contractually covered credit event, 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 fractional loss L^* to the seller of protection in the event of default. While we vary L^* as a parameter, we generally assume that L^* is constant across the sample. To the extent that L^* varies over time or across issuers, our implied risk-neutral default probabilities would be corrupted. This is not crucial, as we shall show, when modeling the CDS rates implied by a given EDF. This robustness also applies to the mark-to-market pricing of old default swaps, which is an increasingly important activity, given that the notional amount of debt covered by default swaps is almost doubling each year, and is expected to reach 4 trillion U.S. dollars in 2004, according to the British Bankers Association (www.bba.org).

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 priority. In any case, the option held by the buyer of protection to deliver from a list of debt instruments will cause the effective fractional loss given default to the seller of protection to be the maximum fractional loss given default of the underlying list of debt instruments. If restructuring is included as a covered default event, the impact of this cheapest-to-deliver option is, within the current "modified-modified" ISDA standard contract, mitigated by a contractual restriction on the types of deliverable debt instruments, especially with respect to maturity.

Ignoring the cheapest-to-deliver effect, the CDS rate is, in frictionless markets, extremely close to the par-coupon credit spread of the same maturity as the default swap, as shown by Duffie (1999). Our results thus speak to the relationship between EDFs and corporate credit spreads. We are told

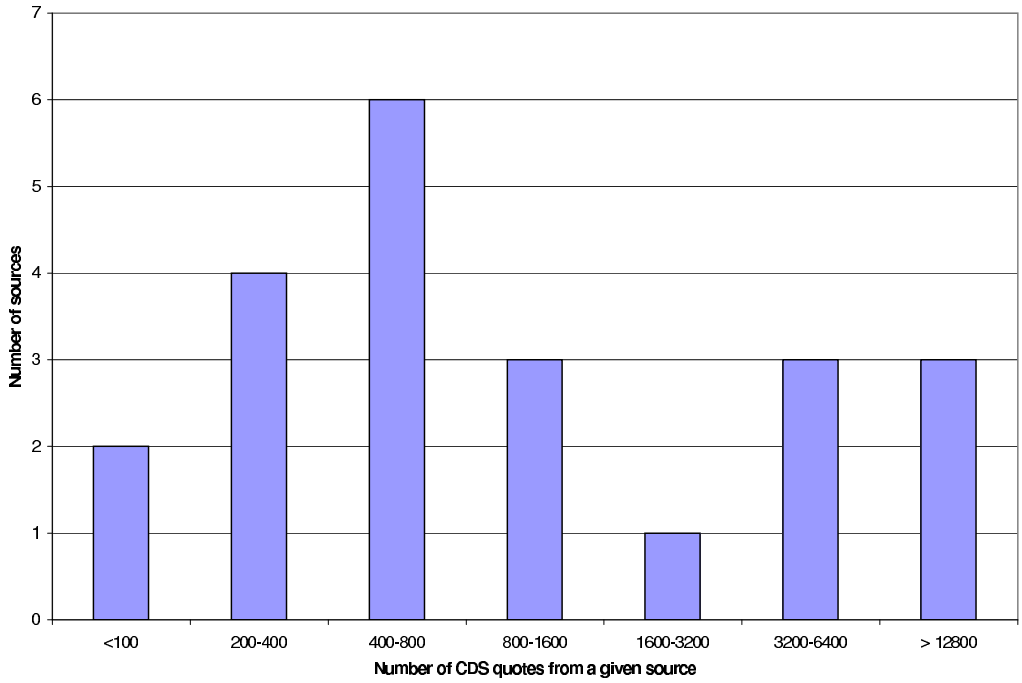

Figure 2: Distribution of CDS quote providers by number of quotes provided

by market participants that asset swaps, synthetic approximations of par-coupon bonds, as explained in Duffie (1999), trade at "par spreads" that are, on average, becoming closer to CDS rates as the CDS market matures and grows in volume, liquidity, and transparency. This is confirmed to some extent in empirical studies by Longstaff, Mithal, and Neis (2003) and Blanco, Brennan, and Marsh (2003), provided one measures bond spreads relative to interest-rate swap yields.

Our CIBC data set consists of over 49,000 intra-day CDS rate quotes on 69 firms from three Moody's industry groups. The sources of these quotes include 12 investment banks and 10 default-swap brokers. The cross-sectional concentration of the number of quotes by source is shown in Figure 2. A breakdown of the number of quotes by banks and by default-swap brokers is given in Table 1.

The three representative Moody's industry groups that we selected for analysis are North American Broadcasting and Entertainment, North Amer-

Table 1: Breakdown of number of CDS quotes by type of source

	Total	Median	Average	Sources	Min	Max
Banks	34462	871	2872	12	116	9523
Brokers	15030	620	1503	10	147	5431
All	49492	702	2250	22	116	9523

ican Oil and Gas, and North American Healthcare. The quotes are all for 5-year, quarterly premium, senior unsecured, US-Dollar-denominated, at-the-money default swaps. A company from any of these 3 sectors is included in our study if and only if at least 50 historical CDS bid/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 3, which shows, for each credit rating, the number of firms in our study of that median Moody's rating during the sample period. Figure 3 indicates a concentration of Baa-rated firms.

Daily CDS mid-point rate quotes were generated from intra-day bid and ask quotes using the following algorithm.

1. If a bid and an ask were present, we record the bid-ask spread.
2. If the bid is missing, we subtract the average bid-ask spread to estimate the ask.
3. If the ask is missing, we add the average bid-ask spread to estimate the bid.
4. From the resulting bid and ask, we calculate the mid-quote as the average of the bid and ask quotes.

The firms that we studied from the broadcasting-and-entertainment industry are listed in Table 2, along with their median 1-year EDF, median Moody's credit rating during the sample period, 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 Tables 8 and 9 of Appendix C, respectively.

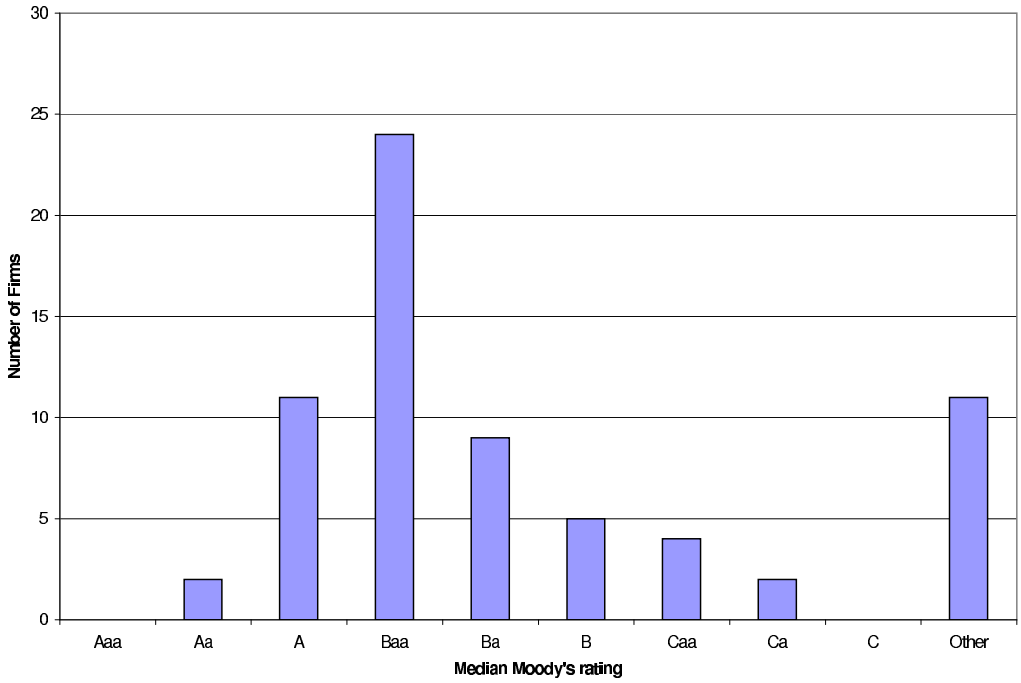

Figure 3: Distribution of firms by median credit rating during the sample period. Source: CIBC.

3 Descriptive EDF-CDS Model

A simple preliminary linear model of the relationship between a firm's 5-year CDS (Y_i) and the 5-year EDF (X_i) measured in basis points on the same day is

$$\begin{aligned} Y_i &= 52.26 + 1.627X_i + e_i, \\ & \quad (1.58) \quad (0.007) \end{aligned} \tag{1}$$

where X_i is the observed 5-year EDF of a given firm on a given day, Y_i is an observed CDS rate of the same firm on the same day, e_i is a random disturbance. Standard errors are shown parenthetically. The ordinary-least-squares coefficient estimates are Based on 18,259 paired EDF-CDS observations from September 2000 to August 2003, with most observations during 2002.

The associated coefficient of determination, R^2, is 0.718. Figure 4 illustrates the fit of (**??**), for all firms in our study, and all time periods. The

Table 2: Broadcasting and Entertainment Firms

Name of Firm	Median EDF (basis points)	Median Rating	No. Quotes
Adelphia Communications	349	B2	279
AOL Time Warner	11	N/A	3447
Charter Communications	281	B3	444
Clear Channel Communications	29	Baa3	1698
Comcast	39	Baa3	1043
COX Communications	36	Baa1	2153
Insight Communications	173	B3	303
Liberty Media	21	Baa3	515
Mediacom Communications	286	Caa1	168
Primedia	65	Ba3	325
Royal Caribbean Cruises	134	Baa2	462
Viacom	13	Baa1	2458
Walt Disney	7	N/A	2745

5-year CDS rate is thus estimated to increase by approximately 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 default intensities (actual and risk-neutral) to be constant, this would imply an average ratio of risk-neutral to actual default intensity φ of approximately (16/0.75)/10, or 2.0.

Linearity of the CDS-EDF relationship, however, is placed in doubt by the sizable intercept estimate of roughly 50 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 quarterly regressions that the associated intercept estimates are unreasonably large in magnitude. We also noted that scatter plots of the CDS-EDF relationship indicated a pronounced concavity at low levels of EDF. That is, the sensitivity of credit spreads to a firm's estimated default probability seems to decline at larger levels of default risk. There is also apparent heteroskedasticity; with greater variance for higher EDFs. The slope of the fit illustrated in Figure 4 is thus heavily influenced by the CDS-to-EDF relationship for

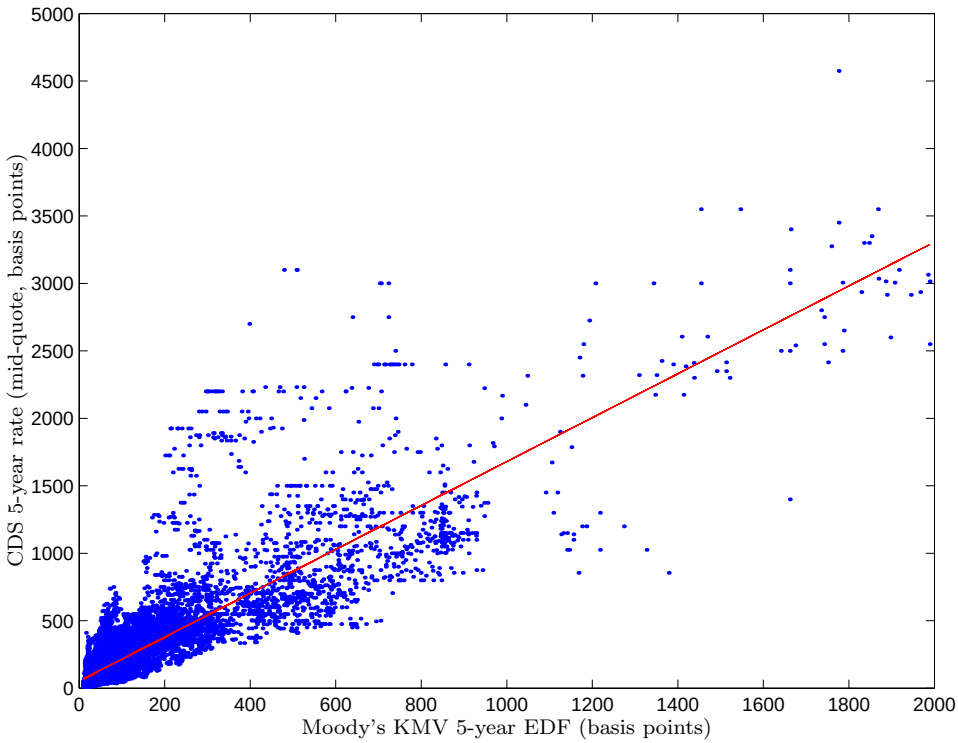

Figure 4: Scatter plot of EDF and CDS observations and OLS fitted relationship. Source: CIBC (CDS) and Moody's KMV (EDF).

lower-quality firms. Accordingly, we next turned to the specification[2]

$$\log Y_i = \alpha + \beta \log X_i + z_i, \tag{2}$$

for coefficients α and β, and a residual z_i. The fit, illustrated in Figure 5, shows much less heteroskedasticity, although some potential corruption from the granularity of EDFs of extremely high-quality firms.

We have taken CDS rate observations (Y_i) by two approaches: *(i)* the daily median CDS for each given name, and *(ii)* all CDS observations for that day. The second approach, which has substantially more CDS observations

[2]We 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.

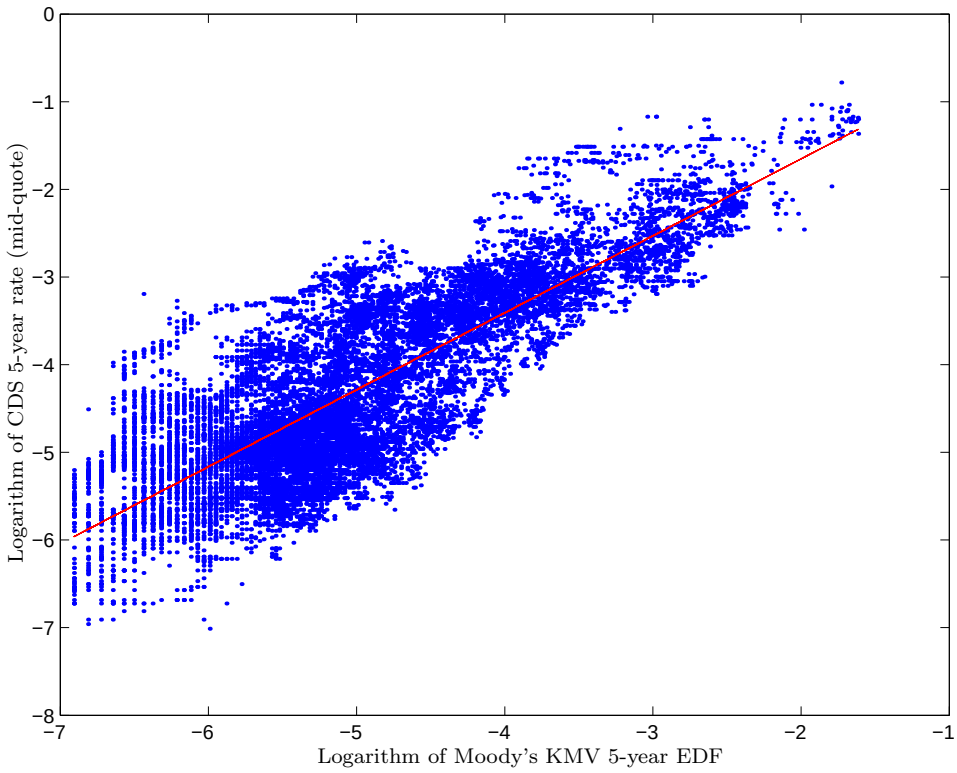

Figure 5: Scatter plot of EDF and CDS observations, logarithmic, and OLS fitted relationship. Source: CIBC (CDS) and Moody's KMV (EDF).

per EDF observation, has by construction a lower coefficient of determination (R^2), and is likely to have more precise estimates of the intercept and slope coefficients, a and b. (This is necessarily so if the model is correctly specified.) It is from this model with more observations that we would thus anticipate getting a more precise notion of how CDS rates are related to EDFs.

Technically, the two cases (daily median CDS observations, and all CDS observations) would not both be consistent with equation (**??**), since the median is an order statistic that depends on sampling noise in a non-linear fashion. We prefer, in any case, the median to the average daily CDS observation as we believe it to be more robust to outliers induced by observation noise. All of our regressions were fit by ordinary least squares.

We anticipate that practitioners will fit CDS rates to EDFs along these

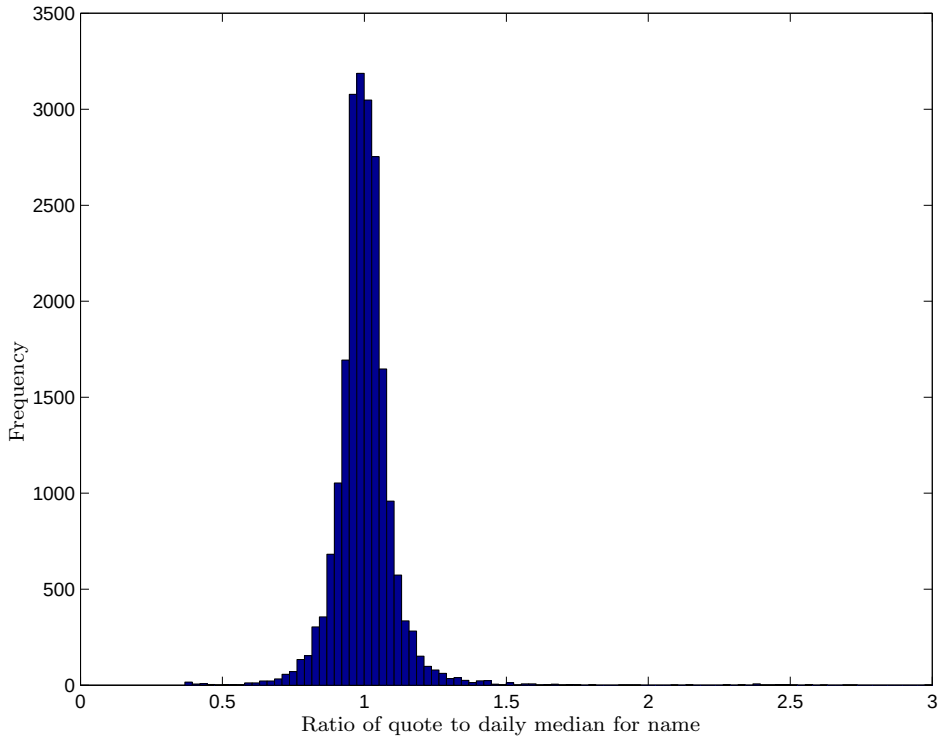

Figure 6: Intraday distribution of ratio of five-year CDS bids to median bid, after removing the median bids. Source: CIBC.

lines a tool for marking-to-market positions for which liquid CDS quotes are not always available.

Figure 6 shows a histogram of the ratio of quotes to the daily median quote for the same name, after removing the points associated with the median quote itself (of which there are approximately 19,500). The plot shows substantial intraday variation in CDS quotes of a given name.

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 highly correlated, presumably because they are both images of the same covariate, distance to default. 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 coeffi-

cients, presumably from multi-collinearity. We do not report the results for the multiple regressions. In any case, the 5-year EDF captures the average effect of default risk over the 5-year period, as does the CDS. This is not to suggest, however, that default risk premia implicit in the CDS rates necessarily have the same term structure. We have little information about this term structure to report at this time. (We plan to later analyze short-maturity CDS data.)

Moving from a pooled regression, we control for changes in the CDS-to-EDF relationship across time and across sectors. Table 10, found in the appendix, 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 ($18,259$ observations in all), including dummy variables for sectors and months. For example, extracting from Table 10 the fit implied for the oil-and-gas sector, we have

$$\log \mathrm{CDS}_i = \underset{(0.029)}{0.912} + \underset{(0.004)}{0.828} \log \mathrm{EDF}_i + \sum \hat{\beta}_j D_{\text{month } j}(i) + z_i, \quad (3)$$

where $\hat{\beta}_j$ denotes the estimate for the dummy multiplier for month j, with j running from December 2000 through August 2003, and z_i denotes the residual. We obtain an R^2 of about 74%. From the one-standard-deviation confidence band implied by normality of the residuals for the logarithmic fit, the associated confidence band for a given CDS rate places it between 58% and 171% of the fitted rate. From the dummy coefficient estimate for the healthcare sector, the CDS rate for a healthcare firm is estimated to be 20% higher than that of an oil-and-gas firm with the same EDF. A broadcasting-and-entertainment firm is estimated to have a 39% higher CDS than an oil-and-gas firm with the same reported EDF. These differences are highly significant, both statistically and in terms of their economic implications for risk premia, or loss given default, or both.

The fitted model also shows highly significant variation in risk premia across the months of 2002 and 2003, with the highest risk premia during the third quarter of 2002, when, for a given EDF, spreads are estimated to have been roughly 50% higher than they were in August, 2003. Figure 7 illustrates this variation over time with a plot of the dummy variables of the regression model 3, indicating the percentage increase in CDS rates at a given EDF assocaited with each month. Figure 8 tells a similar story with a plot of the weekly average of the daily median, by sector, of the ratio of the 5-year CDS

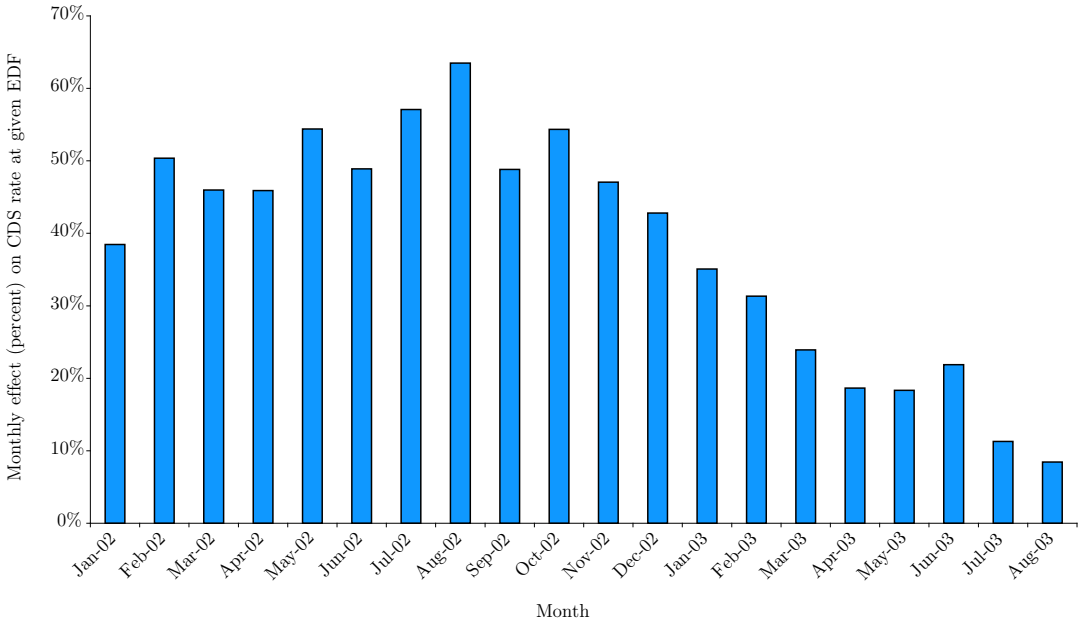

Figure 7: Monthly dummy multipliers in CDS-to-EDF fit.

rate to the 5-year EDF. This index of default risk premium peaks, for every sector, during July and August of 2002.

4 Actual Default Intensity from EDF

The default intensity of an obligor is the instantaneous 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 $(\Omega, \mathcal{F}, P)$ 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

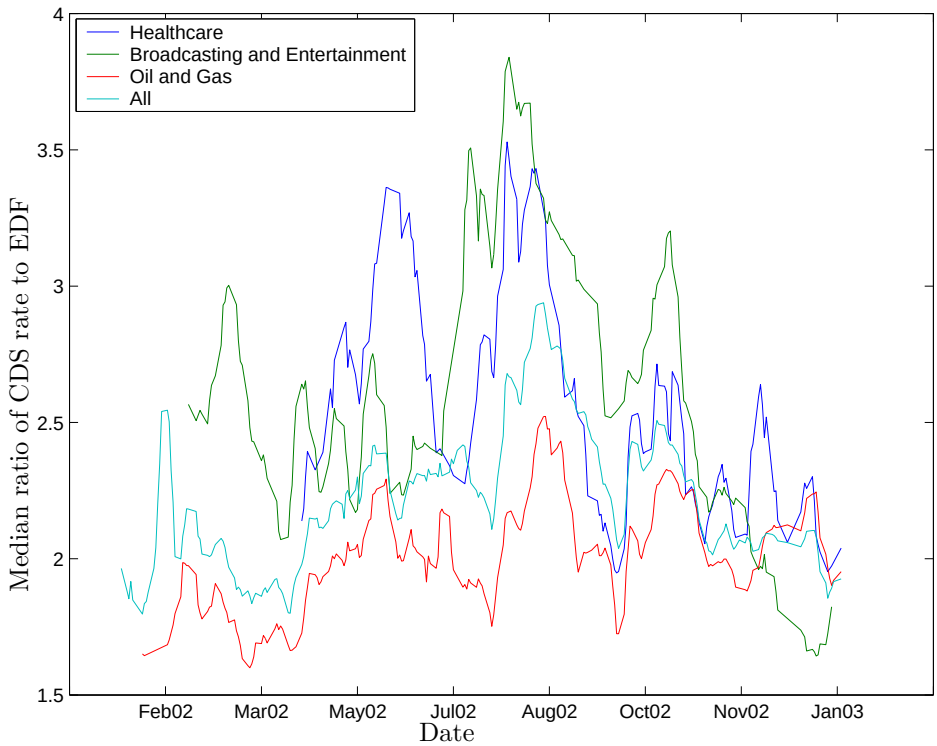

Figure 8: Weekly average of sector-median 5-year CDS-to-EDF ratio.

conditional probability at time t, for a currently surviving obligor, that the obligor survives to some later time T, is

$$p(t,T) = E\left(e^{-\int_t^T \lambda(s)\,ds} \;\middle|\; \mathcal{F}_t\right). \tag{4}$$

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, $t+h$, $t+2h, \ldots$, where h is one month. From these observations, we estimate a time-series model of the underlying intensity process λ, for each firm. In total, we analyzed 69 firms.

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

$$dX_t = (a - \kappa X_t)\, dt + \sigma\, dB_t, \tag{5}$$

where B is a standard Brownian motion, and a, κ, and σ are constants to be estimated. The behavior for $\lambda = e^X$ is sometimes called a Black-Karasinski model.[3] This leaves us with a vector $\Theta = (a, \kappa, \sigma)$ of unknown parameters to estimate from the available monthly EDF observations of a given firm. We have 123 months of 1-year EDF observations for most of the firms in our sample, for the period January, 1993 to August, 2003.

In general, given the log-autoregressive form of the default intensity in (5), there is no closed-form solution available for the 1-year EDF, $1 - p(t, t+1)$ from (4). 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), as well as a more rapid algorithm, explained in the Appendix B, based on approximation of the solution in terms of a basis of Chebyshev polynomials. (Our current parameter estimates are for the trinomial-tree algorithm.)

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

$$\sup_{\Theta} \quad \mathcal{L}\left(\{1 - p(t_i, t_i + 1) : 1 \le i \le N\};\, \Theta\right), \tag{6}$$

where $t_1, t_2, \ldots, t_N$ are the N observation times for the given firm, and $\mathcal{L}$ denotes the likelihood score of observed EDFs 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,

$$G(\lambda(t); \Theta) = 1 - E_\Theta\left(1 - e^{-\int_t^{t+1} \lambda(s)\, ds} \;\middle|\; \lambda(t)\right), \tag{7}$$

   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] See Black and Karasinski (1991).

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, as indicated in Figure 5. We therefore remove from the sample any firm whose sample-mean EDF is below 10 basis points.

4. There were occasional missing data points. These gaps were also treated exactly, assuming the event of censoring is independent of the underlying missing observation.

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

We have not imposed a joint distribution of EDFs across firms. It could be feasible, for example, to impose joint normality of the Brownian motions driving each firm's EDFs, with a specified cross-firm correlation structure. (This is planned for subsequent research.)

Figure 3 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%. Frequency plots of the estimated volatility and mean-reversion coefficients, σ and κ, are shown in Figures 10 and 9, respectively. The estimated parameter vector for each firm is provided in Table 12, found in Appendix C.

One notes significant dispersion across firms in the estimated parameters. Monte-Carlo analysis revealed substantial small-sample bias in the MLE estimators, especially for mean reversion. We therefore obtained sector-by-sector estimates for κ and σ, shown in Table 4, while allowing for a firm-specific drift parameter a, listed in Table 13 in Appendix C.[4]

[4]The intensity λ is measured in basis points.

Table 3: Number of observations of 1-year EDFs. Data: Moody's KMV.

Ticker	uncensored	capped at 20%	total	Ticker	uncensored	capped at 20%	total
ABC	92	0	92	HAL	123	0	123
ABT	123	0	123	HCA	123	0	123
ADELQ	97	16	113	HRC	123	0	123
AGN	123	0	123	HUM	121	0	121
AHC	123	0	123	ICCI	44	0	44
AMGN	123	0	123	JNJ	123	0	123
AOL	123	0	123	KMG	123	0	123
APA	123	0	123	KMI	123	0	123
APC	123	0	123	KMP	121	0	121
BAX	123	0	123	L	81	0	81
BEV	121	2	123	LLY	123	0	123
BHI	123	0	123	MCCC	35	0	35
BJS	123	0	123	MDT	123	0	123
BMY	123	0	123	MGLH	103	19	122
BR	123	0	123	MRO	123	0	123
BSX	123	0	123	NBR	123	0	123
CAH	123	0	123	NEV	123	0	123
CAM	92	0	92	OCR	123	0	123
CCU	123	0	123	OEI	123	0	123
CHIR	123	0	123	OXY	123	0	123
CHK	109	13	122	PDE	123	0	123
CHTR	33	7	40	PKD	123	0	123
CMCSA	123	0	123	PRM	86	3	89
CNG	85	0	85	PXD	123	0	123
COC	45	0	45	RCL	120	0	120
COP	123	0	123	RIG	119	0	119
COX	95	0	95	SBGI	92	0	92
CVX	123	0	123	THC	123	0	123
CYH	76	0	76	TLM	123	0	123
DCX	53	0	53	TRI	46	0	46
DIS	123	0	123	TSO	123	0	123
DO	86	0	86	VIA	123	0	123
DVN	123	0	123	VLO	123	0	123
DYN	0	0	0	VPI	123	0	123
ENRNQ	106	1	107	WFT	123	0	123
EP	122	1	123	WLP	123	0	123
F	123	0	123	WMB	115	8	123
FST	121	0	121	WYE	123	0	123
GENZ	123	0	123	YBTVA	99	0	99
GM	123	0	123				

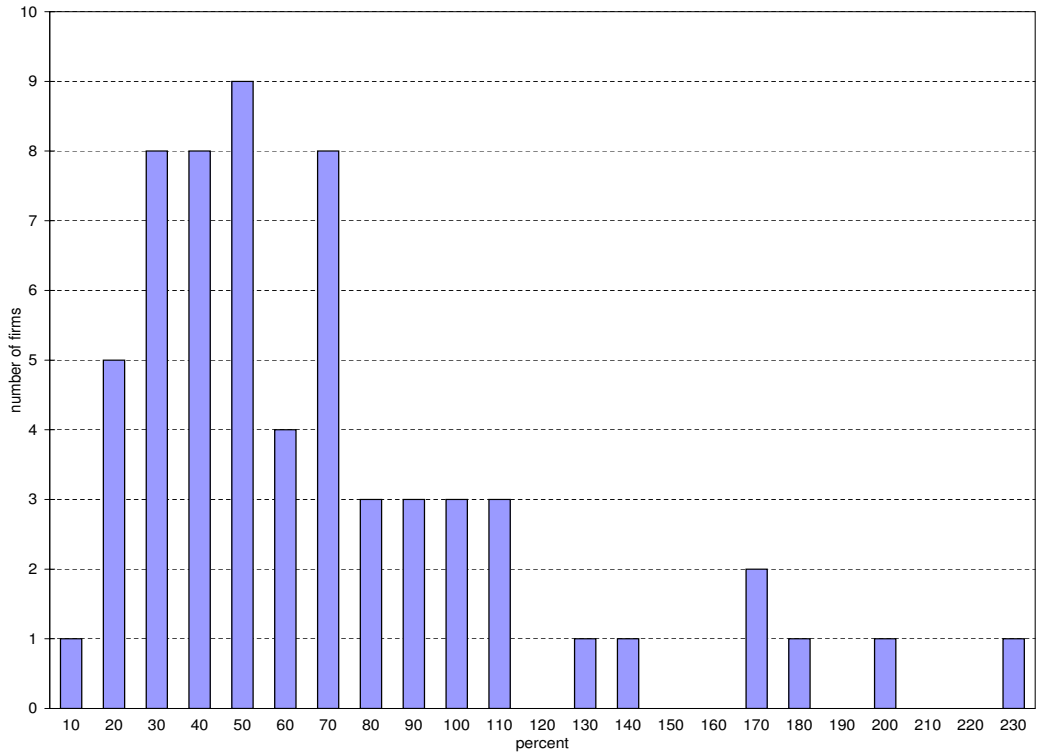

Figure 9: Distribution of estimated default intensity mean-reversion parameters (κ).

5 Risk-Neutral Intensity from CDS and EDF

This section explains our methodology for extracting risk-neutral default intensities, and probabilities, from CDS and EDF data.

5.1 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

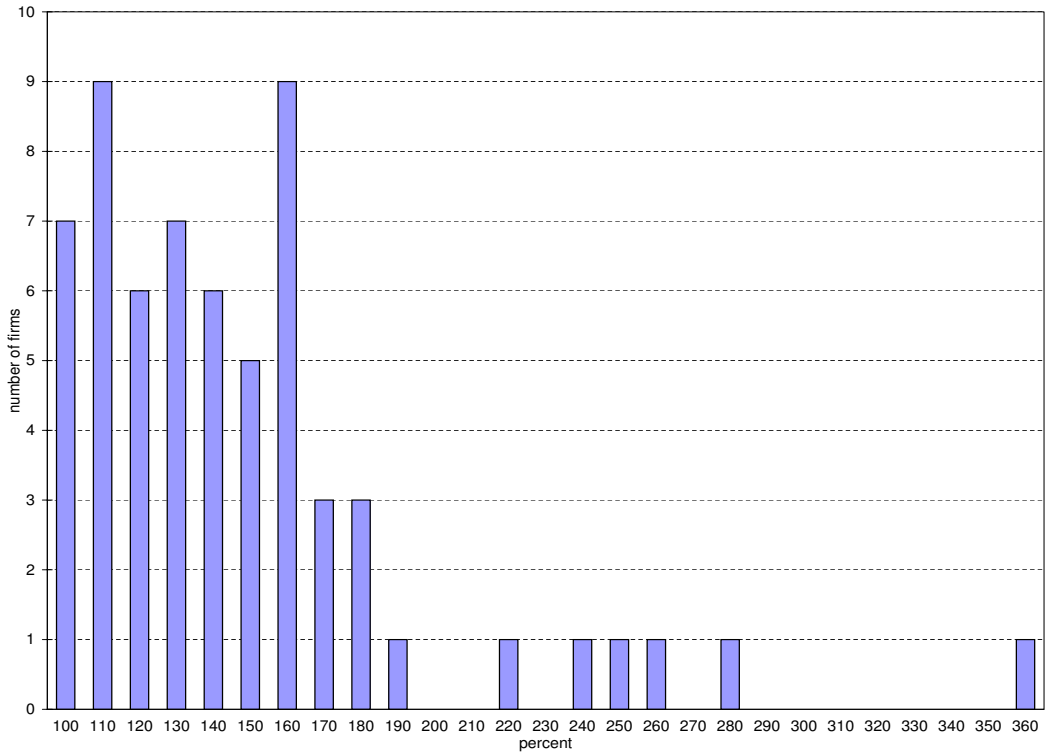

Figure 10: Distribution of estimated default intensity volatility parameters (σ).

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 stopping time $\tau > t$, of

$$S_t = E^Q \left(e^{-\int_t^\tau r(u)\,du} W \;\middle|\; \mathcal{F}_t \right), \tag{8}$$

where r is the short-term interest-rate process[5] and E^Q denotes expectation with respect to an equivalent martingale measure Q, that we fix. One may

[5]Here, 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 $e^{\int_t^T r(s)\,ds}$.

Table 4: Sector EDF-implied default intensity parameters.

	mean($\hat{\theta}$)	$\hat{\kappa}$	$\hat{\sigma}$	no. firms
Oil and Gas	3.3309	0.4663	1.2910	32
Healthcare	3.3784	0.6559	1.5123	17
Broadcasting and Entertainment	4.2625	0.7082	1.6372	14

view (8) 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 double-stochastic property under P, this need not imply the same convenient double-stochastic property under Q as well. Indeed, Kusuoka (1999) gave a counterexample. We will nevertheless assume the double-stochastic property under Q. (Sufficient conditions are given in Duffie (2001), Appendix N.) Thus, we have

$$Q(\tau > T \mid \mathcal{F}_t) = p^*(t,T) = E^Q\left(e^{-\int_t^T \lambda^*(u)\,du} \;\middle|\; \mathcal{F}_t\right), \tag{9}$$

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, Duffee (1998).) It follows from (8) and this independence assumption that the price of a zero-coupon defaultable bond with maturity T and zero recovery at default is given by

$$d(t,T) = \delta(t,T)p^*(t,T), \tag{10}$$

where $\delta(t,T) = E_t^Q\left(e^{-\int_t^T r(s)\,ds}\right)$ is the default-free market discount and $p^*(t,T)$ is the risk-neutral conditional survival probability of (9).

Extensions to the case of correlated interest rates and default times are treated, for example, in Lando (1998).

A default swap stipulates quarterly payments by the buyer of protection of premiums at an annual rate of c, as a fraction of notional, until the default-swap maturity or default, whichever is first. From (10), 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

$$g(t) = \frac{1}{4} \sum_{i=1}^{n} \delta(t, t(i)) p^*(t, t(i)), \tag{11}$$

for premium payment dates $t(1), \ldots, t(n)$. The market value of the potential payment by the seller of protection on this default swap is

$$h(t, c) = E^Q \left(\delta(t, \tau) W_\tau^c 1_{\tau \leq t(n)} \;\middle|\; \mathcal{F}_t \right), \tag{12}$$

for the payment at default, if it occurs at time t, of

$$W_t^c = L_t^* - c \left(t - \frac{\lfloor 4t \rfloor}{4} \right), \tag{13}$$

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.[6] The second term in (13) 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

$$C(t) g(t) = h(t, C(t)). \tag{14}$$

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

$$h(t, c) = \int_t^{t(n)} \delta(t, s) E^Q \left(e^{-\int_t^s \lambda^*(u)\, du} \lambda^*(s) W_s^c \;\middle|\; \mathcal{F}_t \right) ds. \tag{15}$$

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

We take L^* to be constant and use, as a numerical approximation of the integral in (15),

$$h(t,c) \simeq \sum_{i=1}^{n} \delta\left(t, \frac{t(i)+t(i-1)}{2}\right)[p^*(t,t(i-1))-p^*(t,t(i))]\left(L^* - \frac{c}{8}\right), \quad (16)$$

which involves a time discretization of the integral in (15) that, in effect, approximates, between quarter ends, with a linear discount function and risk-neutral survival function. Then $C(t)$ is calculated from (14) using this approximation.

5.2 Model Specification

For a parametric specification of the risk-neutral default intensity process λ^*, motivated by our regression results, we suppose that

$$\log \lambda_t^* = \alpha + \beta \log \lambda_t + u_t, \quad (17)$$

where α and β are constants, $X = \log \lambda$ is as specified earlier by (5), and u satisfies

$$du_t = -\kappa_u u_t \, dt + \sigma_u \, d\xi_t, \quad (18)$$

where, under the actual probability measure P, we take ξ to be a standard Brownian motion independent of the Brownian motion B of (5).

The risk-neutral distribution of (λ^*, λ) is specified by assuming that

$$dB_t = -(\gamma^0 + \gamma^1 X_t)\, dt + d\tilde{B}_t \quad (19)$$

and

$$d\xi_t = -(\gamma_u^0 + \gamma_u^1 u_t)\, dt + d\tilde{\xi}_t, \quad (20)$$

where $(\tilde{B}, \tilde{\xi})$ is a two-dimensional standard Brownian motion under Q, and where $\gamma^0, \gamma^1, \gamma_u^0$, and γ_u^1 are constants. In addition to the parameter vector Θ, the model for λ^* requires an estimator of the parameter vector

$$\Theta^* = (\alpha, \beta, \kappa_u, \sigma_u, \gamma^0, \gamma^1, \gamma_u^0, \gamma_u^1).$$

5.3 Estimation Strategy and Results

For any given firm, 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 parameter vector Θ of the actual intensity model λ following the procedure described in Section 4. In a second step, fixing the estimate of Θ, and treating this estimate as though in fact equal to the true parameter vector, we estimate the parameter vector Θ^* governing the risk-neutral intensity process λ^* on a sector-by-sector basis. For this second step, our data consists of weekly observations of 5-year default swap rates and 1-year EDFs, over a time period from 9/27/00 until 4/9/03. 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 (14), which we approximate using (16). In the current implementation, we only use pairs of CDS-EDF observations where neither the CDS or the EDF data is missing, and for which the EDF is not censored at 20%. In addition, we remove from the sample any firm whose sample-mean EDF is below 10 basis points.

Preliminary investigations have shown that, by restricting Θ^* so that the parameter-implied stationary mean of u is equal to its model-implied sample mean across all firms in a given sector, one obtains considerable improvement in the interpretability of the parameter estimates, facilitating the comparison of the implied values for λ and λ^*. One the same grounds, and because the data were unable to sharply pin down the market-price-of-credit-risk factors γ^0 and γ_u^0 exactly, we set $\gamma_0^u = 0$ and further restrict Θ^* so that the parameter-implied risk-neutral stationary mean of λ^* for a given firm is equal to the sample mean of CDS rates divided by the risk-neutral mean loss rate at default. In addition, we also impose, in this current implementation, that λ^* is a one-dimensional lognormal intensity process under Q, by taking $\tilde{\kappa} = \tilde{\kappa}_u$. Here, $\tilde{\kappa}$ and $\tilde{\kappa}_u$ denote the risk-neutral mean-reversion parameters of X and u, respectively.[7] We anticipate relaxing the latter two restrictions in a subsequent version of the model.

Preliminary sector-by-sector parameter estimates for the broadcasting-and-entertainment, healthcare, and oil-and-gas industries are summarized in Table 5.[8] We allow for a firm-specific market-price-of-default-risk parameter γ^0. Estimates by firm are listed in Table 14, Appendix C. Figure 11 shows

[7] From Equations (5) and (18) through (20) we have $\tilde{\kappa} = \kappa + \gamma^1\sigma$ and $\tilde{\kappa}_u = \kappa_u + \gamma_u^1\sigma_u$.

[8] Both λ and λ^* are measured in basis points.

the implied sample paths of λ^* and λ for Vintage Petroleum, and Figure 12 displays the time series of Vintage's estimated default risk premia, that is, the ratio of its risk-neutral to its actual default intensities.

Table 5: Preliminary sector CDS-implied risk-neutral default intensity parameter estimates

	Oil and Gas	Healthcare	Broadcasting and Entertainment
$\hat{\alpha}$	2.8987	2.4917	4.5272
$\hat{\beta}$	0.4167	0.6290	0.2451
$\hat{\kappa}_u$	1.6218	1.7509	0.8603
$\hat{\sigma}_u$	1.8311	2.2719	1.6911
mean($\hat{\gamma}^0$)	0.4025	0.4637	1.3007
$\hat{\gamma}_1$	−0.0790	−0.0668	−0.2880
$\hat{\gamma}_u^1$	−0.6867	−0.5264	−0.3688
mean(λ^*/λ)	2.53	5.59	5.50
no. firms	29	14	12

For example, extracting from Table 5 the fit implied for the healthcare sector, we have

$$\log \lambda_t^* \quad = \quad 2.49 + 0.63 \log \lambda_t + u_t,$$

or equivalently,

$$\lambda_t^* \quad = \quad 12.06\, \lambda_t^{0.63}\, e^{u_t},$$

where λ_t and λ_t^* are measured in basis points. So, for an actual default intensity of 100 basis points and $u_t = 0$, we get a risk-neutral default intensity of roughly 219 basis points. The risk-neutral distribution of λ^* is estimated as

$$d \log \lambda_t^* = (\hat{\tilde{a}} - 0.55 \log \lambda_t^*)\, dt + 2.46\, dB_t^*,$$

where $\tilde{a} = \hat{a} - \hat{\gamma}^0 \hat{\sigma}$. The sample averages of the estimated risk premia are 2.53, 5.59, and 5.50 for the oil-and-gas, healthcare, and broadcasting-and-entertainment sector, respectively.

As a diagnostic check, we examine the behavior of the standardized innovations $\epsilon_{t+h}, \epsilon_{t+2h}, \ldots$ of u_t, defined by

$$u_{t+h} = e^{-\kappa_u h} u_t + \sigma_u \sqrt{\frac{1 - e^{-2\kappa_u h}}{2\kappa_u}}\, \epsilon_{t+h}.$$

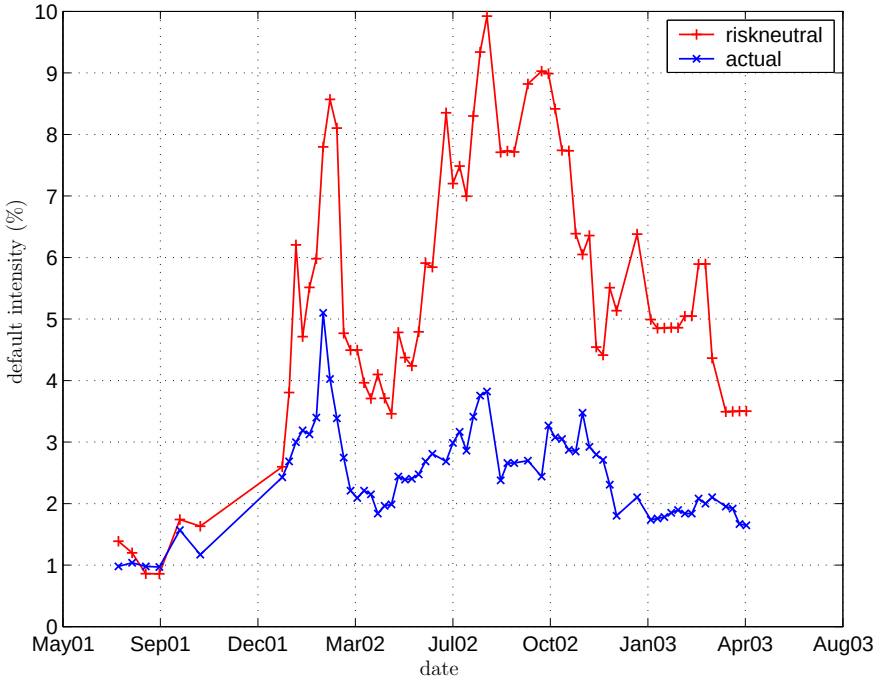

Figure 11: Implied default intensities for Vintage Petroleum. Sources: Moody's KMV and CIBC.

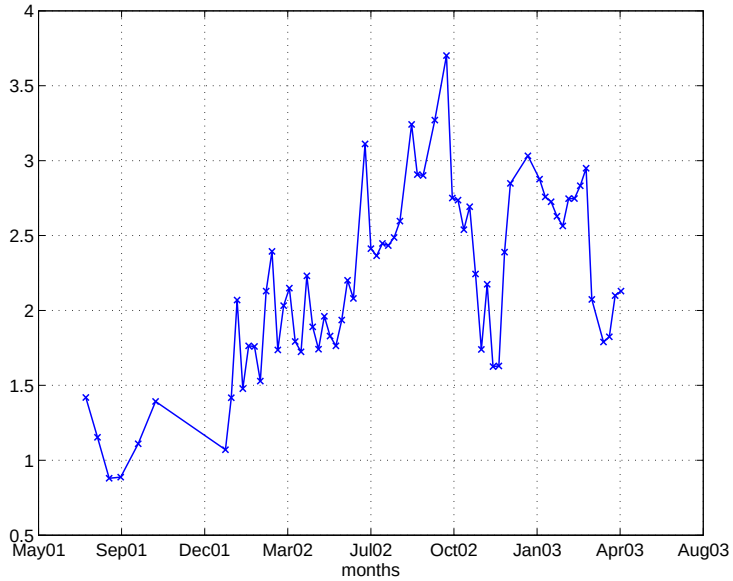

Figure 12: Estimated default risk premia, λ^*/λ, for Vintage Petroleum.

Under the specified model, and under the actual probability measure P, these innovations are independent standard normals. Table 6 lists the sample mean and the sample standard deviation (SD) of the fitted versions of these standardized innovations, for each of the three sectors. Finally, Figure 13 shows the associated histogram of fitted ϵ_t, merging across all firms, plotted along with the standard normal density curve.

Table 6: Sample moments for standardized innovations

	Mean	SD
Healthcare	-0.0101	0.9962
Oil and Gas	-0.0590	0.9760
Broadcasting and Entertainment	-0.0006	1.0008
All	-0.0347	0.9862

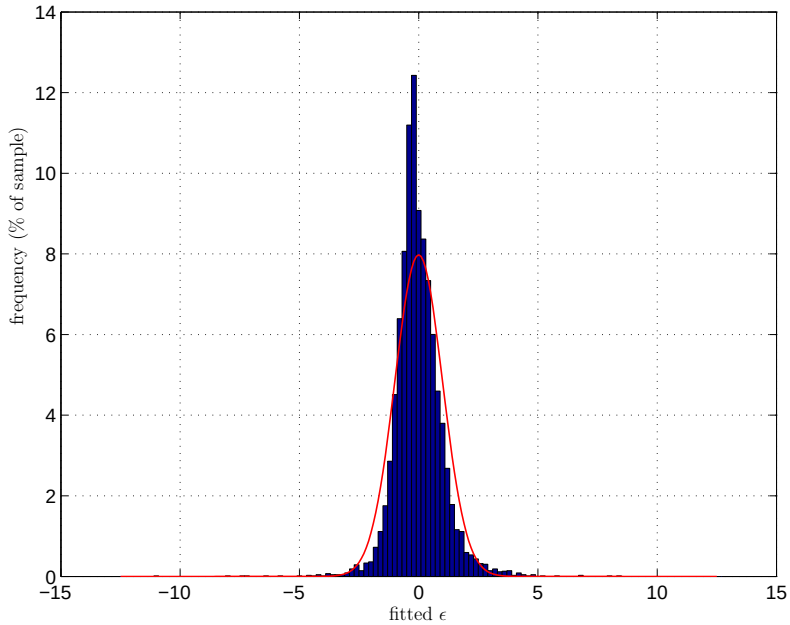

Figure 13: Estimated innovations ϵ across all sectors, and the standard normal density.

6 Discussion and Conclusion

We compare our results on default risk premia to those available in the literature. Using the structural model of Leland and Toft (1996), Huang and Huang (2000) 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, the estimated ratios of annualized risk-neutral to actual five-year default probabilities are reported in Table 7. In magnitude, the results are roughly consistent with those of Driessen (2002). One notes that the risk premium typically declines as default probability increases, as suggested by the results of our basic log-log regression model, and as captured by our time-series formulation.

Table 7: Five-year default risk premium implied by structural-model results of Huang and Huang (2001)

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

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 above 20%. Our data is the monthly observed EDF level Y_i at month i, for each of N month-end times $t_0, t_1, \ldots, t_N$.

From (5), for any time t and time step h (which is 1/12 in our application), the discretely sampled log-intensity process X satisfies

$$X_{t+h} = b_0 + b_1 X_t + \epsilon_{t+h}, \tag{A.1}$$

where $b_1 = e^{-\kappa h}$, $b_0 = a/(\kappa(1-b_1))$, and $\epsilon_{t+h}, \epsilon_{t+2h}, \ldots$ 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 $\Theta = (a, \kappa, \sigma)$ as follows. First, we regress $\log(1-Y_i)$ on $\log(1-Y_{i-1})$, 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 $\Theta = (a, \kappa, \sigma)$ at

$$\begin{aligned}
\kappa_0 &= -\frac{\log(\hat{b}_1)}{h}, \\
a_0 &= \frac{\hat{b}_0}{1-\hat{b}_1}\kappa_0, \\
\sigma_0 &= \hat{\sigma}_\epsilon \sqrt{\frac{2\kappa_0}{1-\exp(-2\kappa_0 h)}}.
\end{aligned}$$

If Θ is the true parameter vector, then $Y_i = G(\lambda(t_i); \Theta)$, where G is defined via (7).

Suppose, to pick an example of a censoring outcome from which the general case can easily be deduced, that for months k through $\bar{k} > k+1$, 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 $l+1$ to month $\bar{l}$ are missing. Let $\mathcal{I} = \{i : k+1 \le i \le \bar{k}\} \cup \{i : l+1 \le i \le \bar{l}\}$ denote the censored month numbers. Then the likelihood of the 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

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

where $P(\,\cdot\,; Y_n = y_n; \Theta)$ denotes the distribution of $\{Y_{n+1}, Y_{n+2}, \ldots\}$ 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

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

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

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

In the remainder of this appendix, we describe how to compute $q(Y;\Theta)$ by Monte Carlo integration, and hence $P(Y_{k+1} \geq \zeta, \ldots, Y_{\bar{k}} \geq \zeta; Y_k = y_k, Y_{\bar{k}+1} = y_{\bar{k}+1}, \Theta)$. 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 $(Z_{k+1}, \ldots, Z_{\bar{k}})$ given Z_k and $Z_{\bar{k}+1}$ which is given by

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

We are now in a position to estimate the quantity in (A.4) by generating some "large" integer number J of independent sample paths $\{(Z_{k+1}^j, \ldots, Z_{\bar{k}}^j); 1 \leq j \leq J\}$ from the joint conditional distribution of $(Z_{k+1}, \ldots, Z_{\bar{k}})$ 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 $\{k+1, \ldots, \bar{k}\}$.

B Solution of Log-Normal Intensity Model

This appendix provides an algorithm, prepared for this project by Gustavo Manso, for computing the survival probability of (4), and related expectations of the form $E(e^{-\int_0^t \lambda(s)\,ds} F(\lambda_t))$, for a well-behaved function $F : [0,\infty) \to \mathbb{R}$. The algorithm allows for a generalization of the log-normal intensity model to a model that is, in logarithms, autoregressive with a mixture-of-normals innovation, allowing for fat tails and skewness. Matlab code is downloadable at the web site www.stanford.edu/∼/manso/numerical/.

Inputs: Parameters $(k, m_1, v_1, p, m_2, v_2, m)$ and initial log-intensity $x \in [a, b]$.

Output: Let $y(j) = \lambda(t_j)$, for equally spaced times $t_0, t_1, \ldots, t_m$. The output is

$$S(0,x) = E\left[\exp\left(-\sum_{j=1}^{m} y(j)\right) F(y(m))\right],$$

where

$$\begin{aligned} \log y(j) &= -k \log y(j-1) + W(j) + Z(j), \\ \log y(0) &= x, \end{aligned}$$

and $W(j)$ is normal, mean m_1, variance v_1, $Z(j)$ is, with probability p, equal to 0 (no jump) and with probability $1-p$, normal with mean m_2, variance v_2. All $W(j)$ and $Z(j)$ are independent.

Step 1 Compute $K \geq N+1$ Chebyshev interpolation nodes on $[-1, 1]$:

$$z_k = -\cos\left(\frac{2k-1}{2K}\pi\right),\ k = 1, \ldots, K.$$

Step 2 Adjust the nodes to the $[a, b]$ interval:

$$x_k = (z_k + 1)\left(\frac{b-a}{2}\right) + a,\ k = 1, \ldots, K.$$

Step 3 Evaluate Chebyshev polynomials:

$$T_n(z_k) = \cos(n \cos^{-1} z_k), \;\; k = 1, \ldots, K \;\; \text{and} \;\; n = 1, \ldots, N.$$

Step 4 Recursive Integration:

- Boundary condition: $S(m, x) = F(\exp(x))$, for $x \in [a, b]$.
- For $j = m : -1 : 0$,

  1. Numerical Integration:

  $$S(j, x_k) = \pi^{-\frac{1}{2}} \sum_{i=1}^{I} \omega_i \left[pq(j+1, u_a(i, x_k)) + (1-p)q(j+1, u_b(i, x_k)) \right],$$

  where

  $$\begin{aligned} q(j, u) &= S(j+1, u) \exp(-\exp(u)), \\ u_a(i, x) &= \sqrt{2v_1}\phi_i + (m_1 - kx), \\ u_b(i, x) &= \sqrt{2(v_1 + v_2)}\phi_i + (m_1 + m_2 - kx), \end{aligned}$$

  and (ω_i, ϕ_i), $i = 1, \ldots, I$, are I-point Gauss-Hermite quadrature weights and nodes.[9]

  2. Compute the Chebyshev coefficients:

  $$c_n = \frac{\sum_{k=1}^{K} S(j, x_k) T_n(z_k)}{\sum_{k=1}^{K} T_n(z_k)^2} \;\; \text{for} \;\; n = 0, \ldots, N,$$

  to arrive at the approximation for $S(j, x)$, $x \in [a, b]$:

  $$\widehat{S}(j, x) = \sum_{n=0}^{N} c_n T_n \left(2\frac{x-a}{b-a} - 1 \right).$$

[9] See Judd (1998), page 262, for a table with (ω_i, ϕ_i).

C Additional Background Statistics

This appendix contains additional background statistics regarding the firms studied. Section 2 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 8: Healthcare firms

Firm Name	Median EDF	Median Rating	No. Quotes
Abbott Laboratories	3	Aa1	235
AmerisourceBergen	200	N/A	311
Amgen	2	A2	776
Baxter International	13	A3	741
Beverly Enterprises	432	B1	256
Boston Scientific	84	Baa3	443
Bristol-Myers Squibb	3	Aaa	504
Cardinal Health	25	A2	323
Chiron	16	Baa1	429
Community Health Systems	173	B2	328
Eli Lilly	5	Aa3	403
Genzyme Corp-Genl Division	14	N/A	242
HCA	46	Ba2	540
Healthsouth	257	Ba1	349
Humana	165	Baa3	393
Medtronic	2	N/A	610
Tenet Healthcare	41	Ba1	1129
Triad Hospitals	183	B1	349
Wellpoint Health Networks	47	Baa2	294
Wyeth	17	A2	698

Table 9: Oil and gas firms

Firm Name	Median EDF	Median Rating	No. Quotes
Amerada Hess	26	Baa1	866
Anadarko Petroleum	19	Baa1	1215
Apache	38	Baa1	688
Baker Hughes	45	A2	738
BJ Services	32	Baa2	245
Burlington Resources	35	A3	590
Chesapeake Energy	279	B3	778
ChevronTexaco	3	N/A	491
Conoco	32	A3	532
ConocoPhillips	15	N/A	1334
Consolidated Natural Gas	9	A2	443
Devon Energy	40	Baa1	1463
Diamond Offshore	46	A3	830
EL Paso	42	Baa2	1307
Enron	12	Baa1	678
Forest Oil	196	B1	397
Halliburton	36	Aa3	430
Kerr-McGee	38	Baa1	596
Kinder Morgan Energy Partners	18	Baa1	632
Kinder Morgan	11	Baa2	505
Marathon Oil	62	Baa2	447
Nabors Industries	27	A3	1224
Occidental Petroleum	67	Baa3	1027
Parker Drilling	323	B1	307
Pride International	145	Ba3	870
Talisman Energy	55	Baa1	226
Tesoro Petroleum	171	Ba3	243
Transocean	32	N/A	1038
Valero Energy	85	Baa3	1131
Vintage Petroleum	118	Ba1	449
Weatherford International	140	Baa1	879
Williams Cos	55	Baa2	749

Table 10: CDS-EDF regression results

Results for Daily Median CDS Data

	Levels Model	Standard Error	Log-Log Model	Standard Error
Number of CDS Samples	18259		18259	
Intercept	-28.877	7.846	0.912	0.029
Slope	1.541	0.007	0.828	0.004
Broadcasting Dummy	104.985	3.187	0.389	0.010
Healthcare Dummy	21.439	3.053	0.196	0.010
Dec-00 Dummy	2.605	29.525	0.545	0.094
Jan-01 Dummy	31.516	22.927	0.554	0.073
Feb-01 Dummy	42.600	21.568	0.594	0.069
Mar-01 Dummy	63.467	21.050	0.704	0.067
Apr-01 Dummy	45.407	18.248	0.556	0.058
May-01 Dummy	35.145	20.701	0.416	0.066
Jun-01 Dummy	22.945	21.569	0.215	0.069
Jul-01 Dummy	18.772	19.937	0.117	0.064
Aug-01 Dummy	-8.447	15.164	0.049	0.048
Sep-01 Dummy	2.341	23.433	0.180	0.075
Oct-01 Dummy	20.670	17.013	0.264	0.054
Nov-01 Dummy	20.341	22.436	0.322	0.072
Dec-01 Dummy	-8.964	69.162	0.194	0.221
Jan-02 Dummy	33.728	12.343	0.385	0.039
Feb-02 Dummy	51.169	11.568	0.504	0.037
Mar-02 Dummy	48.130	9.778	0.460	0.031
Apr-02 Dummy	43.675	9.485	0.459	0.030
May-02 Dummy	63.273	9.417	0.544	0.030
Jun-02 Dummy	55.498	9.631	0.489	0.031
Jul-02 Dummy	130.020	9.265	0.571	0.030
Aug-02 Dummy	152.160	9.375	0.635	0.030
Sep-02 Dummy	101.648	9.617	0.488	0.031
Oct-02 Dummy	124.762	9.230	0.544	0.030
Nov-02 Dummy	105.308	9.413	0.471	0.030
Dec-02 Dummy	74.419	9.884	0.428	0.032
Jan-03 Dummy	52.690	9.594	0.351	0.031
Feb-03 Dummy	43.064	9.892	0.313	0.032
Mar-03 Dummy	19.895	9.626	0.239	0.031
Apr-03 Dummy	15.600	9.618	0.187	0.031
May-03 Dummy	23.330	9.671	0.183	0.031
Jun-03 Dummy	27.770	9.605	0.219	0.031
Jul-03 Dummy	11.421	9.646	0.113	0.031
Aug-03 Dummy	4.165	10.808	0.085	0.034
Sum of Squared Residuals	517328076		5270	
Total Sum of Squares	2064101382		20300	
R^2	0.749		0.740	

Table 11: CDS-EDF regression results

Results for Intraday CDS Data

	Levels Model	Standard Error	Log-Log Model	Standard Error
Number of CDS Samples	40844		40844	
Intercept	-35.010	5.167	1.026	0.023
Slope	1.583	0.005	0.785	0.003
Broadcasting Dummy	58.204	1.552	0.322	0.006
Healthcare Dummy	30.340	1.928	0.229	0.007
Dec-00 Dummy	38.421	23.172	0.618	0.088
Jan-01 Dummy	62.357	17.417	0.655	0.066
Feb-01 Dummy	57.496	16.120	0.633	0.061
Mar-01 Dummy	79.335	15.178	0.789	0.058
Apr-01 Dummy	68.074	12.643	0.631	0.048
May-01 Dummy	51.590	14.461	0.443	0.055
Jun-01 Dummy	37.881	15.176	0.263	0.058
Jul-01 Dummy	25.792	13.116	0.194	0.050
Aug-01 Dummy	5.740	10.373	0.125	0.040
Sep-01 Dummy	5.060	16.124	0.207	0.061
Oct-01 Dummy	35.723	11.398	0.313	0.043
Nov-01 Dummy	36.374	15.471	0.391	0.059
Dec-01 Dummy	18.565	57.107	0.278	0.218
Jan-02 Dummy	45.780	7.755	0.484	0.030
Feb-02 Dummy	59.139	7.147	0.592	0.027
Mar-02 Dummy	49.396	6.082	0.531	0.023
Apr-02 Dummy	33.833	5.833	0.456	0.022
May-02 Dummy	61.333	5.873	0.575	0.022
Jun-02 Dummy	43.272	5.883	0.509	0.022
Jul-02 Dummy	112.068	5.708	0.622	0.022
Aug-02 Dummy	128.001	5.781	0.674	0.022
Sep-02 Dummy	91.658	5.854	0.534	0.022
Oct-02 Dummy	101.796	5.674	0.558	0.022
Nov-02 Dummy	89.819	5.775	0.499	0.022
Dec-02 Dummy	63.924	6.096	0.442	0.023
Jan-03 Dummy	42.351	6.127	0.355	0.023
Feb-03 Dummy	36.320	6.168	0.341	0.024
Mar-03 Dummy	21.240	6.078	0.276	0.023
Apr-03 Dummy	20.482	6.109	0.250	0.023
May-03 Dummy	21.402	6.093	0.221	0.023
Jun-03 Dummy	25.479	5.989	0.239	0.023
Jul-03 Dummy	9.514	6.196	0.117	0.024
Aug-03 Dummy	8.619	6.823	0.103	0.026
Sum of Squared Residuals	792797345		11500	
Total Sum of Squares	2952660064		35702	
R^2	0.731		0.678	

Table 12: Fitted parameters of default intensity models

Ticker	$\hat{a}$	$\hat{\kappa}$	$\hat{\sigma}$	Ticker	$\hat{a}$	$\hat{\kappa}$	$\hat{\sigma}$
ABT	†	†	†	MGLH	1.09	0.16	1.09
ADELQ	1.17	0.18	1.40	MRO	1.20	0.45	0.95
AGN	0.83	0.48	1.01	NBR	3.18	1.06	1.62
AHC	1.28	0.58	1.14	NEV	1.34	0.26	0.97
AMGN	†	†	†	OCR	0.69	0.24	1.32
AOL	0.90	0.20	1.10	OEI	1.38	0.33	1.23
APA	4.13	1.38	1.53	OXY	1.94	0.71	0.92
APC	0.93	0.32	0.89	PDE	3.98	0.89	1.70
BAX	1.56	0.64	1.19	PKD	0.79	0.11	1.22
BEV	1.11	0.17	1.14	PXD	1.87	0.46	1.38
BHI	1.24	0.44	0.88	THC	1.43	0.45	1.01
BJS	2.14	0.70	1.35	TLM	0.94	0.38	1.28
BMY	†	†	†	TSO	2.26	0.53	1.33
BR	0.89	0.42	1.04	VIA	1.41	0.64	1.52
BSX	1.12	0.57	1.75	VLO	0.89	0.28	1.04
CAH	1.04	0.49	1.19	VPI	3.20	0.76	1.67
CCU	0.97	0.38	1.57	WFT	0.77	0.24	1.20
CHIR	2.31	0.85	1.24	WLP	1.93	0.68	1.47
CMCSA	2.04	0.56	1.03	WMB	‡	‡	‡
CNG	†	†	†	WYE	†	†	†
COP	†	†	†	CHK	3.42	0.67	1.75
CVX	†	†	†	HUM	1.73	0.37	1.45
CYH	3.89	0.90	1.61	RCL	0.95	0.26	1.14
DIS	†	†	†	RIG	0.81	0.28	1.43
DVN	0.92	0.34	1.46	DYN	‖	‖	‖
ENRNQ	‡	‡	‡	YBTVA	3.66	0.67	1.44
EP	‡	‡	‡	COX	1.79	0.67	1.51
F	0.79	0.23	1.01	ABC	8.17	1.99	2.75
FST	4.79	1.05	1.53	CAM	3.22	0.94	1.32
GENZ	1.87	0.79	1.52	SBGI	3.27	0.61	1.54
GM	3.56	1.09	1.29	PRM	4.69	1.29	2.47
HAL	1.23	0.41	1.53	DO	0.75	0.31	1.54
HCA	2.07	0.87	2.56	L	0.54	0.05	0.95
HRC	0.70	0.14	1.26	DCX	2.22	0.49	1.32
JNJ	†	†	†	COC	4.81	1.78	1.79
KMG	0.46	0.13	0.86	TRI	8.20	1.66	1.22
KMI	3.28	2.20	3.57	ICCI	5.36	0.96	2.14
KMP	1.01	0.37	1.09	CHTR	‡	‡	‡
LLY	†	†	†	MCCC	9.02	1.62	2.30
MDT	†	†	†				

† 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 13: Fitted parameters of default intensity models

Ticker	$\hat{a}$	Ticker	$\hat{a}$	Ticker	$\hat{a}$	Ticker	$\hat{a}$
ABC	2.8120	CHK	2.2229	HRC	2.3440	PKD	2.0882
ABT	†	CHTR	5.0986	HUM	2.6410	PRM	2.9253
ADELQ	3.4884	CMCSA	2.7299	ICCI	4.1459	PXD	1.8506
AGN	1.3468	CNG	†	JNJ	†	RCL	2.4718
AHC	1.1826	COC	1.4114	KMG	1.4174	RIG	1.1622
AMGN	†	COP	†	KMI	0.8738	SBGI	3.7737
AOL	2.4591	COX	2.0626	KMP	1.2905	THC	2.2631
APA	1.4619	CVX	†	L	1.4229	TLM	1.2036
APC	1.3539	CYH	2.9018	LLY	†	TRI	3.3708
BAX	1.5443	DIS	†	MCCC	4.4058	TSO	2.0046
BEV	3.2271	DO	1.0159	MDT	†	VIA	1.7943
BHI	1.4137	DVN	1.2629	MGLH	3.7885	VLO	1.5457
BJS	1.4458	DYN	‡	MRO	1.3825	VPI	1.9995
BMY	†	ENRNQ	1.2787	NBR	1.4524	WFT	1.6014
BR	1.2160	EP	2.0104	NEV	2.3467	WLP	1.9974
BSX	1.2483	FST	2.1402	OCR	1.8883	WMB	1.4543
CAH	1.6181	GENZ	1.5829	OEI	1.8829	WYE	†
CCU	1.5381	HAL	1.2546	OXY	1.3733	YBTVA	3.9456
CHIR	2.0109	HCA	1.0847	PDE	2.1021		

† No estimates provided, mean EDF is less than 10 basis points.

‡ Firm removed from data set.

Table 14: Fitted parameters of risk-neutral default intensity models

Oil and Gas		Healthcare		Broadcasting and Entertainment	
Ticker	$\hat{\gamma}^0$	Ticker	$\hat{\gamma}^0$	Ticker	$\hat{\gamma}^0$
AHC	0.2853	ABC	0.7332	ADELQ	1.6181
APA	0.6787	BAX	0.3906	AOL	1.1280
APC	0.5199	BEV	0.6026	CCU	0.5715
BHI	0.6298	BSX	0.1852	CHTR	2.4714
BJS	0.4423	CAH	0.5710	CMCSA	1.2592
BR	0.3957	CHIR	0.7550	COX	0.8936
CHK	0.6863	CYH	0.6022	ICCI	1.9274
COP	0.5617	GENZ	-0.2647	L	0.4696
DO	0.2172	HCA	-0.3030	MCCC	2.1241
DVN	0.2591	HRC	0.1678	PRM	1.2443
ENRNQ	0.1907	HUM	0.8434	RCL	1.0038
EP	0.3221	THC	0.4496	VIA	0.8981
FST	0.6808	TRI	1.0802		
HAL	-0.0905	WLP	0.6793		
KMG	0.4686				
KMI	-0.0102				
KMP	0.3446				
MRO	0.4744				
NBR	0.5577				
OXY	0.4792				
PDE	0.5742				
PKD	0.5145				
RIG	0.2383				
TLM	0.3290				
TSO	0.4562				
VLO	0.3695				
VPI	0.4796				
WFT	0.6459				
WMB	-0.0289				

References

Altman, E., B. Brady, A. Resti, and A. Sironi (2003). The Link between Default and Recovery Rates: Theory, Empirical Evidence and Implications. 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.

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.

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 (2003). An Emprical Analysis of the Dynamic Relationship between Investment Grade Bonds and Default Swaps. Bank of England.

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.

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.

Driessen, J. (2002). Is Default Event Risk Priced in Corporate Bonds? Working Paper, University of Amsterdam.

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. and K. Singleton (2003). *Credit Risk.* Princeton University Press.

Duffie, D. and K. J. Singleton (1999). Modeling Term Structures of Defaultable Bonds. *Review of Financial Studies* **12**, 687-720.

Duffie, D. and K. Wang (2003). Forecasting Long-Horizon Default with Stochastic Covariates. Working Paper, Graduate School of Business, Stanford University.

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.

G. Delianedis Geske, R. 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.

Harrison, M. and D. Kreps (1979). Martingales and Arbitrage in Multiperiod Securities Markets. *Journal of Economic Theory* **20**, 381–408.

Huang, J. and M. Huang (2000). How Much of the Corporate-Treasury Yield Spread is Due to Credit Risk?: Results from a New Calibration Approach. Working Paper, Graduate School of Business, Stanford University.

Hull, J. and A. White (1994). Numerical Procedures for Implementing Term Structure Models I: Single Factor Models. *Journal of Derivatives* **2**, 7-16.

Judd, K. (1998). *Numerical Methods in Economics.* London, England: The MIT Press.

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 (2000). Analyzing Rating Transitions and Rating Drift with Continuous Observations. Working Paper, Department of Statistics, University of Copenhagen.

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 (2003). Corporate Yield Spreads: Default Risk or Liquidity, Evidence from the Default Swap Market. University of California, Los Angeles.

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.

Shumway, T. (2001). Forecasting Bankruptcy More Accurately: A Simple Hazard Model. *Journal of Business* **74**, 101–124.