

**MODELING
METHODOLOGY**
FROM MOODY'S KMV**CDS-implied EDFTM Credit Measures
and Fair-value Spreads**

Authors

Douglas Dwyer
Zan Li
Shisheng Qu
Heather Russell
Jing Zhang

Contact Us

Americas
+1-212-553-5160
clientservices@moodys.com

Europe
+44.20.7772.5454
clientservices.emea@moodys.com

Asia (Excluding Japan)
+85 2 2916 1121
clientservices.asia@moodys.com

Japan
+81 3 5408 4100
clientservices.japan@moodys.com

Abstract

In this paper, we present a framework that links two commonly used risk metrics: default probabilities and credit spreads. This framework provides credit default swap-implied (CDS-implied) EDF (Expected Default Frequency) credit measures that can be compared directly with equity-based EDF credit measures. The model also provides equity-based Fair-value CDS spreads (FVS) that can be compared directly with observed CDS spreads.

CDS-implied EDF credit measures and fair-value spreads are powerful tools that risk managers can use to extend coverage of credit risk measures, enhance the assessment of default risk, and assess the relative value of various credits. With CDS-implied EDF credit measures, we can provide default risk measures for the population of entities without traded equity, such as private firms, subsidiaries of public firms, and sovereigns, based on their CDS.

For firms with both EDF credit measures and CDS-implied EDF credit measures, risk managers can use both metrics to enhance their assessments of credit risk at the entity level. That is, by comparing information from both markets in a common metric and understanding the differences, risk managers can gain valuable insights into the credit risk of these entities. By using both measures, they can minimize the model risk of relying on one measure alone and increase predictive power of credit risk measures. Additionally, fair-value spreads can be used for mark-to-market valuation and as well as for portfolio management.

Table of Contents

1	Overview.....	5
2	Practical Applications of CDS-implied EDF Measures and Fair-value Spreads	6
2.1	Expanding Coverage of Risk Assessment and Valuation	6
2.2	Practical Applications of EDF and CDS-implied EDF	7
2.3	Practical Applications of the FVS Framework.....	9
2.4	Triangulating Information from Different Sources	9
3	CDS-implied EDF Model and Fair-value Spread Model.....	15
3.2	The Model Framework	16
3.3	Deriving Risk-neutral Default Probability	17
3.4	Bridging Risk-neutral and Physical Default Probability	18
3.5	CDS-implied EDF Credit Measures and Fair-value Spreads	18
3.6	Model Estimation	19
3.7	Model Validation	21
4	Applying the Framework to Sovereigns.....	26
4.1	Implementing the Framework.....	28
4.2	Validation	29
5	Summary	35
Appendix A	Frequently Asked Questions.....	36

1 Overview

Default risk is the uncertainty surrounding an obligor's ability to fulfill its contractual obligations. Defaults are rare events with highly adverse consequences. Thus, assessing default risk is one of a risk manager's most important tasks. In assessing default risk, a prudent risk manager should use all relevant information, which may come from different sources. For corporate credits, the financial statements and general business environment information usually serve as the starting point for fundamental analysis, which typically involves some subjective judgment. Additionally, forward-looking information is widely available in the financial markets, and quantitative approaches can be used to extract default risk signals from market information.

Large amounts of market information exist in both the equity market and the credit market. Both a firm's equity and debt can be valued as options on the asset value of the firm. This insight, originally introduced by Black, Scholes, and Merton, can be applied to assess the default risk of firms with traded equity.¹ In this framework, default occurs when the value of the firm's assets is insufficient to allow the firm to meet its contractual obligations. The unobservable business value of the firm is inferred from equity prices, together with the company's capital structure, and the business risk of the firm. Over the past two decades, Moody's KMV has developed the Vasicek-Kealhofer (VK) version of this framework to calculate Expected Default Frequency (EDF) credit measures for publicly traded companies.² This model, widely used by institutions around the globe, provides daily default probabilities for more than 30,000 public companies worldwide and has been proved to be a forward-looking measurement of default risk for publicly traded firms.³

Another reflection of the default risk of firms with traded debts, such as bonds and loans, can be found in the prices, or equivalently, the yields of these traded debt instruments. The credit spread (i.e., the differences between the yield of the debt instrument and an equivalent default-risk free treasury bond) contains information on the likelihood of default of the borrowers. However, the credit spreads of bonds and loans also reflect other factors, such as recovery risk, market risk premium, embedded options in such instruments, and other non-credit components such as liquidity and taxes. The lack of depth and liquidity of corporate bond and loans markets, relative to the equity markets, also poses a significant challenge to extracting default risk signals for these names.

Another source of credit spreads is the CDS market. A CDS contract is a derivative instrument protecting against default risk, with the buyer paying the seller a premium in exchange for the recovery of credit loss when default happens. The premium (i.e., the CDS spread) mostly reflects default risk—if the likelihood of default is high, the protection is more expensive, which is reflected in a higher spread. Intuitively, credit spreads reflect expected loss, after accounting for investors' risk aversion. Our framework relates credit spreads to their drivers: default probability, loss given default, and market risk premium. With this framework, we can derive a CDS-implied EDF credit measure as a measure of default probability, as well as a fair-value CDS spread (FVS) implied by an equity-based EDF measure. A fair-value CDS spread is an estimate of the CDS spread, calculated from the equity-based EDF credit measure.

By providing both CDS-implied EDF credit measures and fair-value CDS spreads, we can place EDF measures and spreads on a level playing field and integrate signals of credit risk from both the credit and equity markets. If default probabilities are the preferred measure of risk, CDS-implied EDF values can be used in conjunction with EDF credit measures. If spreads are the preferred risk measure, fair-value CDS spreads can be used with observed CDS spreads. This framework helps risk managers extend coverage of credit risk measures derived from market information and helps investors assess relative values of various credits. For example, with CDS-implied EDF credit measures, we can provide default probability measures for the population of entities that do not have EDF credit measures, but have active CDS contracts, including private firms, subsidiaries of public firms, and sovereigns. For firms with both EDF measures extracted from the equity market and CDS-implied EDF from the CDS market, there are now two

¹ Black and Scholes, 1973; and Merton, 1974.

² See Kealhofer, 2003a.

³ Dwyer and Korablev, 2007; Korablev and Qu, 2009.

measures of default risk expressed in a common metric. By using both measures, risk managers can minimize the model risk of relying on one measure alone and potentially increase the predictive power of default detection.

In this paper, we discuss the conceptual framework and estimation methodology of the CDS-implied EDF and FVS model. We also illustrate the applications of both measures in credit risk management and portfolio management and provide validation results showing the performance of these models.

The remainder of the paper is organized in the following way.

- Section 2 describes the practical applications of CDS-implied EDF and FVS.
- Section 3 describes the modeling framework, how the parameters are estimated for corporate credits, and how the framework is validated.
- Section 4 documents the extension of the model to sovereigns.
- Section 5 summarizes the paper and provides concluding remarks.
- Appendix A presents frequently asked questions.

2 Practical Applications of CDS-implied EDF Measures and Fair-value Spreads

CDS-implied EDF and FVS extend the coverage of markets-based metrics for the measurement of default risk, and help investors assess relative values of various credits. We discuss practical applications of the CDS-implied EDF and FVS framework in this section.

2.1 Expanding Coverage of Risk Assessment and Valuation

Although the EDF model provides extensive default risk assessment coverage for publicly traded companies worldwide, many entities with CDS spreads do not have EDF credit measures. Examples include private companies, subsidiaries of public firms, and sovereigns. The default risk of these entities can be assessed with CDS-implied EDF credit measures.

Similarly, using spreads alone to quantify risk in a credit portfolio poses the challenge that many entities in the portfolio may not have spreads from liquid credit instruments. The FVS framework allows one to derive a fair-value spread from an EDF, which can be used to extend spread coverage to entities where spreads are unavailable or unreliable.

2.1.1 CDS Coverage vs. Equity-based EDF Coverage

Table 1 Coverage of CDS-implied EDF Credit Measures on February 22, 2010

	CDS-implied EDF based on a spread provided by at least one dealer quote*	Covered in MIR**
Public Firms	1464	946
Private Firms	258	97
Subsidiaries of Public Firms	473	324
Subsidiaries of Private Firms	218	100
Sovereigns	84	69
State-owned & Supra-national	57	43
Municipalities	88	13

*The minimal requirement for producing a CDS-implied EDF is a spread provided by one dealer.

** Moody's Market Implied Ratings (MIR) provides a CDS implied rating for names that are rated by Moody's Investors Service and the pricing information is viewed as reliable.

EDF coverage by number of entities is substantially greater than for CDS spreads, as many more names have listed equity than traded debt with which to write the CDS contract against. Nevertheless, there are CDS spreads for 1,464 public firms with EDF credit measures. For these names, we can now produce both a CDS-implied EDF credit measure and an EDF credit measure. There are also CDS spreads for 473 subsidiaries of public firms. Examples of these include Merrill Lynch and Wachovia. We define a subsidiary as an entity whose "Moody's Issuer Num" differs from that of its ultimate parent. In many cases, the credit risk of the subsidiary is comparable to that of the parent. There are 258 private firms, which are names with a quoted CDS spread that we are unable to map to an EDF credit measure. Examples of private names in this list include Levi-Strauss and Cargill. There are 218 subsidiaries of private firms, including Residential Capital (a subsidiary of GMAC) and Manor Care (a subsidiary of HCR Healthcare, which is owned by the Carlyle Group). There are 84 sovereigns, 94 state owned enterprises and supra-nationals and 88 municipalities.

To get a sense of EDF coverage for entities with CDS spreads, we can look at the constituents of the North American HY CDX index (CDX.NA.HY.13-v2). The 100 names in this index are chosen by dealers on the basis of CDS contract liquidity. In this index, 13 were private companies, eight were subsidiaries of private companies, and 11 were subsidiaries of public companies. Hence, these 32 entities do not have EDF credit measures due to a lack of traded equity.

2.2 Practical Applications of EDF and CDS-implied EDF

Default risk models, such as the framework for EDF and CDS-implied EDF, can be used by risk managers, lenders, and portfolio managers in a wide range of applications, expanded upon below.

Early Warning and Monitoring

The credit quality of obligors can change quickly, as history has repeatedly demonstrated. Because they are objective and forward-looking, EDF and CDS-implied EDF measures can help risk managers target their resources for risk assessment and mitigation toward cases where they can be the most effective. As shown in Section 2.4, taking the worst of the two measures yields a better credit measure (in terms of Cumulative Accuracy Profile) than taking either of the two individually.

Inputs to Internal Risk Rating Systems

Internal rating systems serve as the foundation of many business decisions within financial institutions: credit approval, limit setting, regulatory compliance, risk based pricing, and active portfolio management. Many institutions have found that market information, when available, is particularly relevant and powerful in internal risk rating assessment. EDF and CDS-implied EDF credit measures are market-based measures constructed to reflect all the relevant information regarding default risk from their respective markets. Thus, they are particularly useful in serving as the market assessment component of an internal rating system.

Benchmarking and Calibrating Internal Risk Rating Systems

An internal rating system needs to provide sufficient differentiation of default risk. To calibrate such a system, regulators typically expect a sizeable amount of realized default events data, spanning at least a full economic cycle. Many institutions may not possess enough internal data for this purpose and can benefit from external sources. EDF credit measures and CDS-implied EDF credit measures are well suited for benchmarking and calibrating internal risk systems, since their development is based on actual defaults, and their usefulness has been demonstrated through extensive validation.

Input to Regulatory Compliance

The probability of default associated with an internal rating plays a central role in the calculation of capital requirements in the Basel II framework. The regulatory capital formula in Basel II uses the so-called Asymptotic Single-Risk Factor (ASRF) and builds on the probability of default. Many banks use external PD models as part of their internal ratings, either for regulatory capital calculations or for benchmarking and calibrating their internal models during their process of fulfilling regulatory requirements.⁴

Input to Required Economic Capital Calculation

Required economic capital (EC) is a concept that financial institutions use to measure their portfolio risk. It can be thought of as the amount of capital the financial institution needs to hold in order to ensure withstanding losses within a given time horizon, consistent with a given solvency probability. The solvency probability is typically derived from the institution's senior debt rating. Computing EC usually requires a portfolio model with probability of default as a key input. A simulation-based portfolio model builds the distribution of credit losses from Monte Carlo simulations of possible portfolio outcomes. Required EC corresponds to the value of the loss distribution at the target probability of solvency, say 99.7%. In general, most of the widely used portfolio models rely upon physical probability of default, as in the case of the Moody's Analytics Portfolio Manager™ and RiskFrontier®.

Loss Provisioning

Loan loss provisions are expenses charged to a bank's earnings when adding to the allowance for possible bad debt. In estimating the provisioning amount, one can use a default risk model to estimate the likelihood of borrowers defaulting on their loans. The model should respond to changes in the risk environment across the economy as a whole. In other words, a provisioning calculation should be as forward-looking as possible. In fact, both International Accounting Standard Board (IASB) and Basel Committee on Bank Supervision are moving toward the more forward-looking "expected loss" approach from the practice of the "incurred loss" approach (e.g., Basel Committee on Banking Supervision, 2009). Both EDF credit measures and CDS-implied EDF credit measures are forward-looking default risk measures that respond to changes in the credit cycle and produce accurate estimates of credit losses over a long period. Consequently, they are appropriate for expected loss-based provisioning calculations.

⁴ Please see BIS document, Use of Vendor Products in the Basel II IRB Framework http://www.fma.gov.at/cms/basel2/attachments/4/8/2/CH0329/CMS1144314485035/bcbs_n18.pdf.

2.3 Practical Applications of the FVS Framework

The idea behind the FVS framework is that a “fair” value credit spread should reflect the underlying risk drivers of the credit, such as default risk and loss given default. The fair-value spreads are constructed to match observed CDS spreads on average. The FVS can be very useful because it is a modeled spread where the model incorporates equity prices. For names where spread information is not readily available, the FVS can be used for valuation purposes. Also, when an entity’s FVS differs from its actual CDS spread, this view can potentially lead to trading opportunities that exploit the relative price difference. These applications are expanded upon below.

Mark-to-market Valuation

Mark-to-market, or fair-value accounting, refers to the accounting standard of assigning a value to a position held in a financial instrument based on the current fair market price for the instrument or similar instruments. Statements of Financial Accounting Standards No. 157, Fair-value Measurements (commonly known as FAS 157), establishes a hierarchy of valuation methodologies. The hierarchy gives first priority to using actual prices for identical assets in active markets, when available, for establishing a fair market value (Level 1 inputs). Second priority is given to valuation methodologies based on inputs that include a combination of prices from inactive markets on identical assets, prices of similar assets from active markets combined with observable characteristics of the asset, and market-corroborated inputs (Level 2 inputs). The lowest priority is given to unobservable inputs, including the firm’s own assumptions regarding how the market would view a particular asset were it to trade.

Fair-value spread produced by our framework can be viewed as Level 2 inputs to a valuation methodology. The fair-value spread takes information from a liquid market (equity prices) and creates an estimate of what the spread on debt would be were the debt to trade actively, using the characteristics of the debt and aggregate information on comparable firms. This application can be used for the thousands of firms that have liquid equity prices but illiquid debt. Examples of illiquid debt include debt that never trades, debt that will trade in the future, and debt that has stopped trading.

A fair-value spread can be used to estimate a benchmark price on such debt were it to trade. A firm may issue a bond to repay a bank debt, in which case its debt may start trading. The fair-value spread can also be used to estimate what the debt will trade at when it starts to trade. The market for a specific firm’s debt may become illiquid for a variety of reasons and, as a result, the current market price for the debt may not be observable. Finally, a fair-value spread can be used to provide an estimate of what the price of the debt should be.

Active Portfolio Management

Portfolio management entails making decisions about taking on additional exposures, selling or hedging existing exposures, and the prices at which to do so. You may also be able to negotiate the terms of a loan in a way that changes the Loss Given Default (LGD) expectation, such as receiving additional collateral on a loan or a “covenant-lite” loan. You can compute the FVS under different LGD assumptions to see how the terms of a loan should impact the loan’s price. Such a framework is key for risk managers to make informed decisions regarding which loans to make, under what terms, and at what price. In addition, the gap between equity-based FVS and observed CDS spreads can be used as a measure of relative mispricing between the equity and CDS market. Risk managers can use this gap to construct investment strategies that exploit the relative price differences between the markets.⁵

2.4 Triangulating Information from Different Sources

For a given set of exposures, information regarding default risk may be available from one or more sources. For example:

- The company may be publicly listed and have readily observable stock prices, in which case, an EDF measure can be computed.

⁵ See Li and Zhang, 2010.

- The company may have liquid CDS transactions and readily available CDS spreads.
- The company may have a published rating from a rating agency.

Although risk managers may prefer to use all available information, it can be difficult to compare, combine, or cross-check signals. For example, when a company has an EDF of 0.5%, an agency rating of Baa2, and a CDS spread of 100 basis points (bps), are these three measures consistent with one another? If not, which measure is inconsistent with the others? The ability to answer such questions facilitates deeper research that could lead to valuable insights on a firm's credit risk.

With our CDS-implied EDF and FVS frameworks, these different information sources can be converted into comparable measures. For example, if the probability of default is the preferred common metric, PDs can be derived from either CDS spreads (as CDS-implied EDF credit measures) or ratings (as rating implied PDs). Since the average EDF level for firms with a given spread level varies considerably over time, as well as by region, sector, and rating, the mapping from CDS to CDS-implied EDF is dynamic and incorporates these factors. In December 2009, a CDS spread of 100 bps for a typical investment grade North American company translated to a CDS-implied EDF of approximately 2%, whereas in April 2008, it translated to an CDS-implied EDF of approximately 0.5%. A rating can be translated into a PD either by computing historical default rates by rating grade or by calculating the average EDF level by rating. The latter approach produces PDs more comparable to EDF and CDS-implied EDF credit measures.

If spreads are the preferred metric, EDF credit measures can be translated into spreads as EDF Implied Fair-value Spreads (FVS). Ratings can be translated into spreads based on median spread by rating. Since our FVS model is dynamic and incorporates region, sector, and rating, we can capture the variation in spread levels over time and across region, rating, and sector. For example, in December 2009, the FVS for an average North American Investment Grade company and EDF of approximately 1% was around 140 bps, in contrast to around 70 bps for an average Investment Grade Japanese company with a similar EDF level.

For a rating scale analysis metric, EDF and CDS spreads can be converted into Equity Implied and CDS implied ratings, as done in Moody's Market Implied Rating tool.⁶

2.4.1 Information Sources

After the information sources are converted to comparable measures, these measures will often differ. To understand the potential drivers of these differences, it is important to understand the unique characteristics of these information sources and their impacts on default risk assessment. There are a number of possible explanations for the differences between the EDF and CDS-implied EDF for certain credits, explained as follows.

Parent and Subsidiaries

Common stock can be issued by the holding company of a corporate family, while its subsidiaries can have different debt issues. In fact, the debt holders may prefer for the debt to be issued by the operating company because, in the event of default, their claim on the assets of the operating company would typically be senior to debt issued by the holding company. Debt issued by the holding company is termed to be *structurally subordinate* to debt issued by the operating company. CDS contracts can be written with reference to either debt issued by the holding company or those by a specific subsidiary. In these cases, the EDF credit measures pertain to the corporate family as a whole, while the CDS-implied EDF credit measures may more accurately reflect the default risk of the issuing subsidiary.

For example, United Utilities Water PLC is a regulated utility and the main operating subsidiary of United Utilities PLC. Debt issued by the subsidiary can be safer than debt issued by the parent, because of what is termed *structural subordination*. This difference is reflected in their different CDS spreads and shown in Figure 1. The difference in the CDS spreads is then reflected in the differences in the CDS-implied EDF credit measures. The equity-based EDF credit measure, however, does not reveal such a difference because the equity-based EDF credit measure pertains to

⁶ See Munves et al., 2007.

the corporate family as a whole. In this case, using CDS-implied EDF would provide insights into the different risk levels within the corporate family that the EDF credit measure cannot.

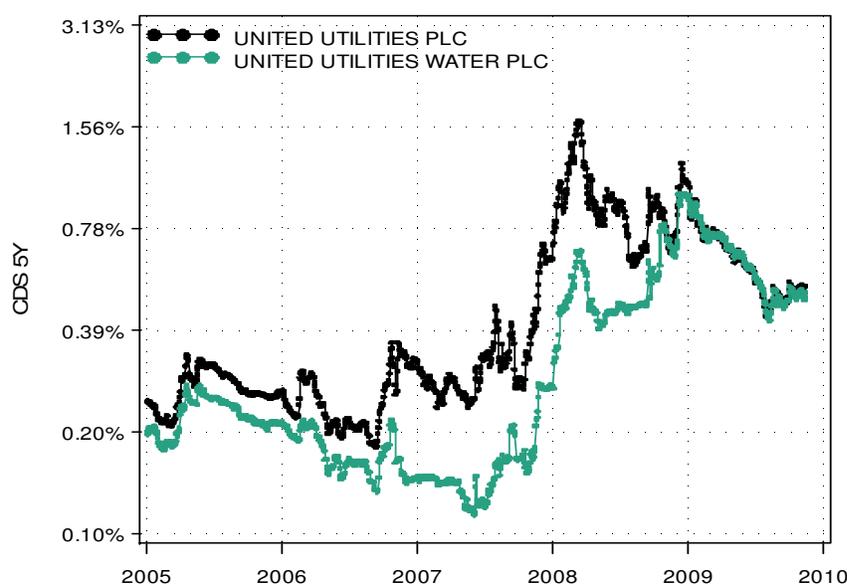


Figure 1 Spreads of Utilities: Holding and Operating Company

Significant Change in Capital Structure

The EDF model translates the market price of equity into a measure of default risk using the market value of assets, the volatility of this value, and the liability structure of the firm. The EDF credit measure has been extensively validated and has been used widely by market participants to manage credit risk. Nevertheless, it is a model, and how the model is implemented is important for the signals of risk that it provides. While great efforts have been taken to ensure the latest financial statement information is used in calculating the EDF credit measure, there are circumstances when the most recent capital structure change may not be immediately reflected in the EDF measures. In such cases, it is more than likely that CDS spreads have already incorporated the information. These capital structure changes include issuing a large amount of new debt, mergers and acquisitions, divestures, and leveraged buy-outs (LBO).

For example, following the announcement of an LBO, the equity-based EDF credit measure typically will not reflect the increase in credit risk due to the increase in expected leverage associated with deal.⁷ The LBO of Hospital Corporation of America (HCA), the largest private operator of health care facilities in the world, illustrates this scenario. On July 24, 2006, HCA agreed to sell itself to three private-equity firms. The transaction was to be financed by the firm issuing more debt to buy back their own stock. When the deal was announced, the credit market immediately recognized the greater default risk for HCA due to the increased leverage, and the firm's CDS spread almost tripled overnight. On the other hand, the reported EDF measure, not capturing the change in capital structure, showed a decrease in risk due to increased share prices, as shown in Figure 2.

⁷ The EDF credit measures reported in our product are based on published financials statements. During an LBO, merger, acquisition, or spin-off, there is often good reason to believe that the published financials do not reflect what the capital structure of the post-deal entity will be. In such circumstances, we recommend clients use Moody's Analytics EDFCalc™ to compute the EDF on the basis of an estimate of what the capital structure will be after the completed deal.

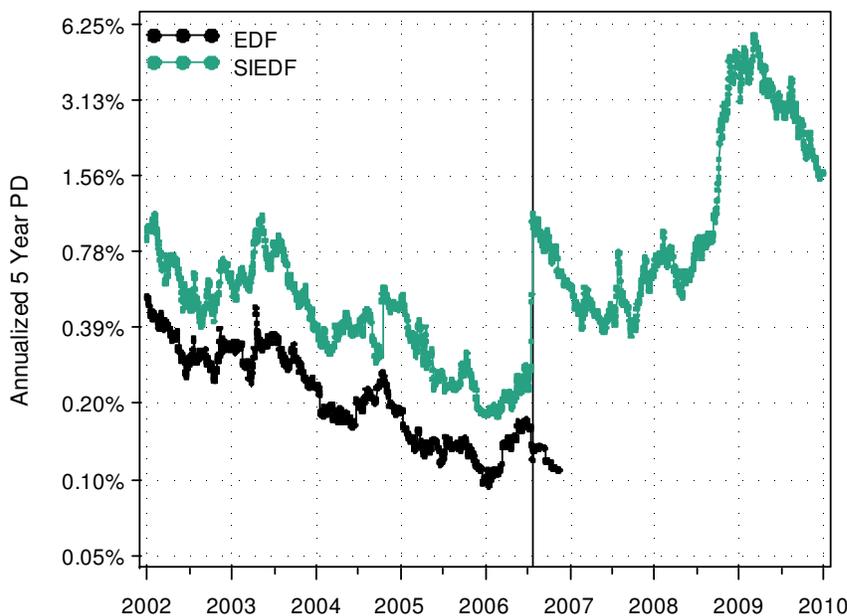


Figure 2 EDF and CDS-implied EDF of HCA Inc.

Equity and CDS Markets May Reflect Different Concerns

The EDF measure is designed to measure the default risk of the entire capital structure of a firm, whereas CDS contracts are usually written with reference to a specific debt instrument, to establish the seniorities of deliverable securities. Sometimes these debt instruments have explicit or implicit government guarantees, causing spreads to be low relative to EDF credit measures. For example, years before being placed into conservatorship by the U.S. Treasury, Fannie Mae and Freddie Mac had EDF measures much higher than implied by both their CDS spreads. The CDS contracts referenced bonds that had an implicit guarantee from the U.S. government and therefore traded at a relative low spread.

Since an EDF measures the probability of any default in the entire capital structure, a much higher EDF should have raised caution regarding securities further down the capital structure (e.g., preferred stock).⁸ In this case, the comparison between equity and credit is made easier by using default probability as the common risk measure. Many holders of Freddie Mac and Fannie Mae preferred stock may have been lulled into a false sense of security implied by the very low credit spreads on the institutions' senior obligations. Unfortunately, the government support did not apply to the preferred stock, so there was substantial credit risk in this portion of the capital structure that the CDS spreads did not capture.

⁸ One can potentially use the difference between the spread and the spread implied by the EDF (the FVS) to estimate the value of the implied government guarantee (Gray and Malone, 2008).

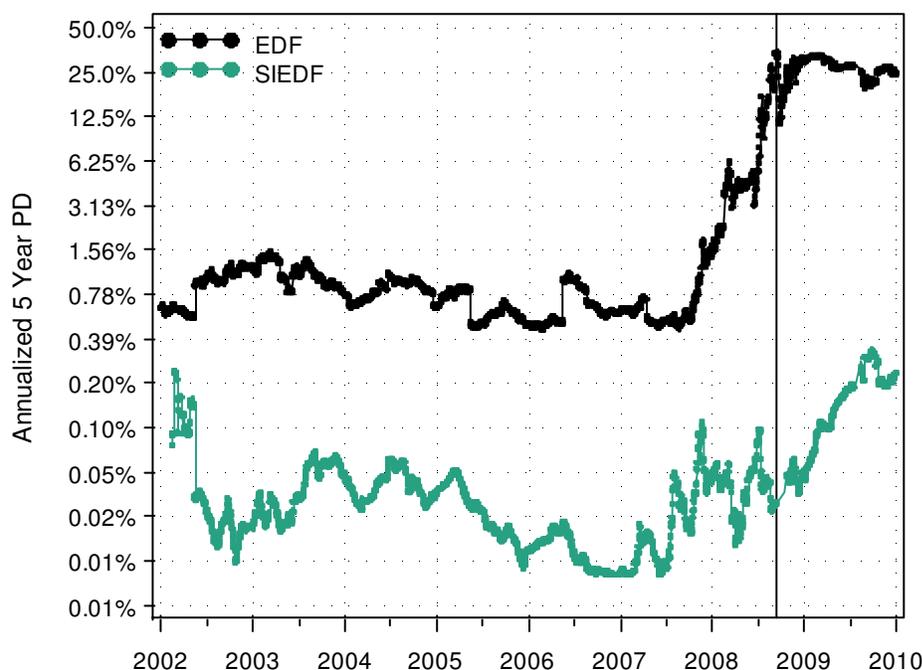


Figure 3 5-year EDF vs. 5-year CDS-implied EDF for Freddie Mac

Signals from the CDS market can differ from signals from equity-based EDF credit measures for many reasons. Understanding these differences can help avoid being lulled into a false sense of security—as may have happened to holders of Fannie Mae and Freddie Mac preferred stock—or misunderstanding the credit implications of a leveraged transaction such as the LBO of HCA.

While understanding these differences can provide valuable insights into risk assessment, it may not be possible to fully attribute them to known factors or to make a judgment on which measure is more relevant than the other. In these cases, we recommend using whichever measure is more conservative (i.e., assigns a higher default probability). This recommendation follows the principle of conservatism in risk management—in the absence of a good reason to discard one risk measure in favor of another, the more pessimistic assessment is used. This approach is particularly relevant for default risk given the asymmetric pay-off of credit. Unlike investing in equity, it does not pay to be optimistic in investing in credit, since the upside on buying a bond is typically bounded.⁹ This recommendation is further supported by empirical evidence that the conservative measure tends to outperform other combinations of the two measures in predicting default, when using Accuracy Ratios to measure performance.¹⁰

For example, between 2001 and 2008, 2,715 firms had both EDF credit measures and CDS spreads, and there were 72 unique default events among these firms. The predictive power, as measured by the Accuracy Ratio, of the conservative measure (i.e., the maximum of EDF and CDS-implied EDF) is 84.6%, measurably higher than 77.3% of EDF credit measures and 79.4% of CDS-implied EDF credit measures. Figure 4 reports power curves.

⁹ In certain circumstances, a bond investor can earn large returns. For example, a bond investor may earn a large return when they invest in distressed debt that then recovers or invest in long dated bonds when interest rates fall. Nevertheless, in most cases the distribution of bond returns will be skewed to the left.

¹⁰ We have tested linear combinations of the logarithms of the PDs, the PDs themselves, and the normal inverse of the PDs.

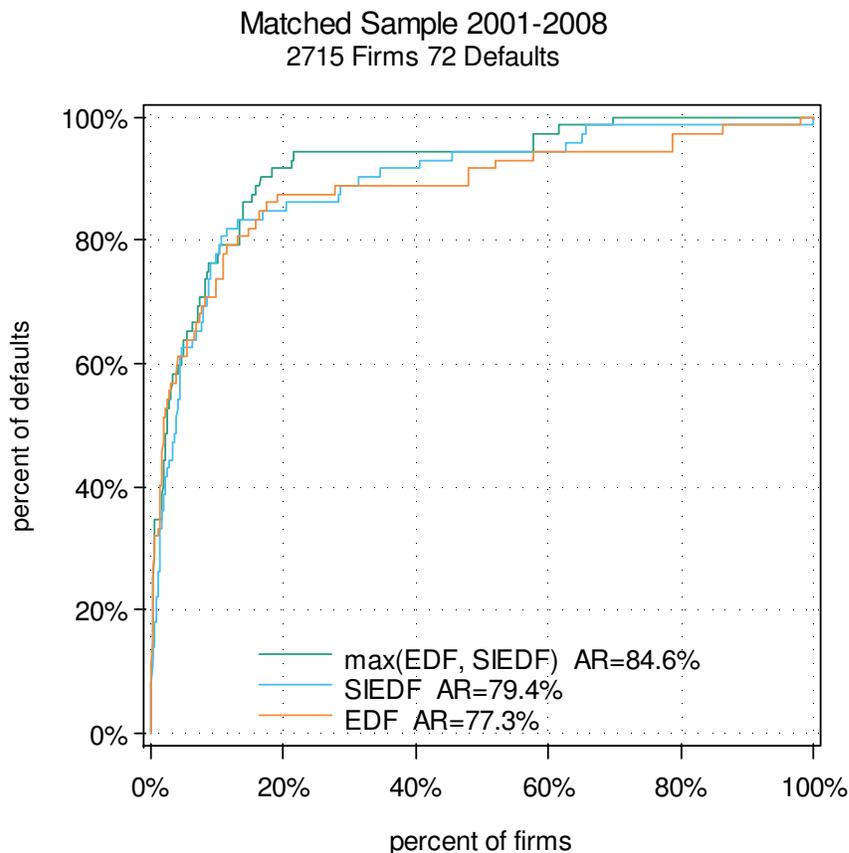


Figure 4 Power Curves of EDF Credit Measures and CDS-implied EDF Credit Measures on the Matched Sample

In a typical bank portfolio, most obligors do not have a CDS-implied EDF measure as most exposures do not trade in the CDS market. For such portfolios, we can still make a conservative PD measure utilizing both EDF and CDS-implied EDF credit measures: the conservative measure for a firm is the same as its EDF, if a CDS-implied EDF is not available and is the maximum of EDF and CDS-implied EDF when both exist.

We have tested this approach on a group of large public companies. Our criteria for large firms were chosen to minimize the problem of missing defaults (i.e., defaults that occurred but not found in our database).¹¹ Between 2001 and 2008, 20,849 such unique companies existed, with 737 unique default events. On this sample, the Accuracy Ratio by the conservative measure is 81.3%, higher than 80.6% of EDF credit measures alone and 80.1% of EDF credit measures substituted by CDS-implied EDF when available, as seen in Figure 5. Since only about 13% of these firms have CDS-implied EDF credit measures, the higher accuracy ratio for the conservative measure suggests considerable improvement in default detection for this portion of the sample. The advantage of the conservative measure appears to be small, but this is because most companies do not have CDS-implied EDF, and for them, the conservative measure is the same as the EDF.

¹¹ We include non-financial firms with sales greater than \$100 million at some point during their history and financial firms with book assets greater than \$1 billion at some point during their history.

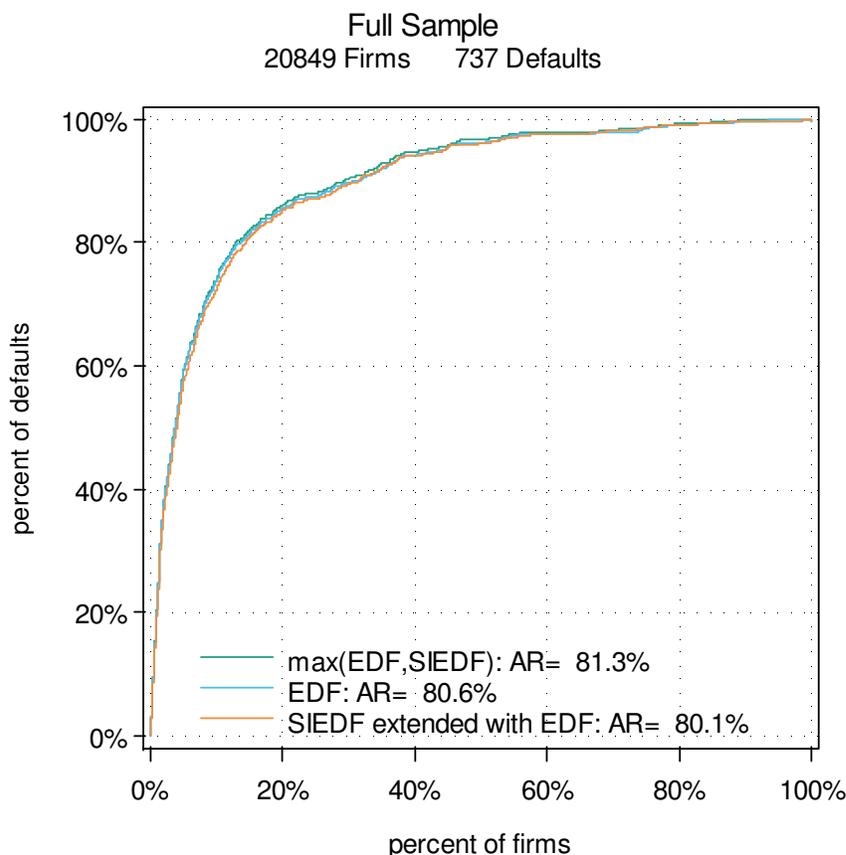


Figure 5 Power Curves of EDF Credit Measures and CDS-implied EDF Credit Measures on EDF Sample

In summary, this framework can generate a CDS-implied EDF from a CDS spread and a Fair-value Spread from an EDF. It provides distinct and comparable views on the risk of a credit from two different markets. The differences in these views are valuable. When one market is signaling a higher level of risk than the other, using the more conservative signal leads to a better Accuracy Ratio than using either one by itself or by using the CDS-implied EDF when it is available and the EDF otherwise.

3 CDS-implied EDF Model and Fair-value Spread Model

The CDS-implied EDF Model translates credit spreads into comparable physical default probabilities as measured by EDF credit measures. We calibrate the model so that predicted and observed CDS spread levels are comparable for each region, asset class, and sector. The model is invertible in the sense that the same theoretical framework and parameters that we use to derive CDS-implied EDF credit measures from CDS spreads can also be used to derive fair-value spreads from EDF credit measures. The model framework is based on applying the theory of risk-neutral pricing to the valuation of debt instruments subject to default risk.

3.1.1 Risk-neutral Pricing and Risk-neutral Probabilities of Default

Suppose that all market participants were risk-neutral. Then the task of valuing a risky asset would simply involve calculating the expected cash flows using actuarial fair probabilities of default and then discounting them to the present at the risk-free rate. Determining these probabilities might be difficult, but otherwise, the calculation is straightforward.

In reality, investors are risk-averse; market participants prefer a certain payoff to a variable one with the same expected value. To compensate for the uncertainty, investors require a premium for the embedded credit risk. To value instruments in a manner consistent with market prices, this risk premium must be taken into account during the valuation process. Consequently, the future contingent cash flows need to be adjusted for the risk involved in actually receiving them.

The idea of risk-neutral pricing is to adjust probabilities to account for risk and then discount expected cash flows at the risk free rate, rather than computing expected values directly and then discounting at a risk-adjusted discount rate. Those adjusted, “virtual” probabilities are called risk-neutral probabilities. They constitute the risk-neutral measure.¹²

These risk-neutral probabilities are set so that the price of a security is equal to the present value of the expected future cash flows, computed with the adjusted probabilities rather than the actual probabilities. In the context of instruments with default risk, the risk-neutral probability of the default state of the world has come to be known as the risk-neutral probability of default.¹³

3.2 The Model Framework

The model framework includes two basic components. The first component links CDS spreads with risk-neutral default probabilities. The second component links risk-neutral default probabilities to physical default probabilities. We calibrate the model using observed CDS spreads and equity-based EDF measures so that the model can be used to convert spreads to equivalent EDF levels and, in the opposite direction, to convert EDF levels to fair-value CDS spreads. Both directions involve the intermediate step of finding an equivalent term structure of risk-neutral default probabilities, as depicted in Figure 6.

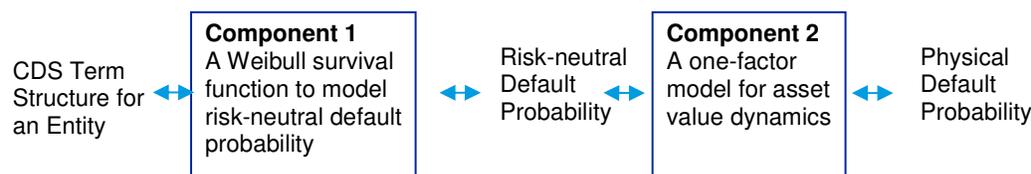


Figure 6 Model Components

Component 1

The first model component translates between CDS spreads and risk-neutral default probabilities. Recall, given the risk-neutral probability of receiving a contingent cash flow, the present value of the cash flow is its expected value, computed using the risk-neutral probability of receiving it, discounted to the present using the risk-free rate. For a CDS contract, from a contract buyer’s perspective, future cash flows include premium payments and the recovery of the credit loss in the event of default. As these payments are contingent on default, their present values depend on the risk-neutral default term structure.

The present value of the default payment also depends on the expected LGD. Under the risk-neutral probabilities, the default protection price (i.e., the present value of the premium payments) should be equal to the present value of the expected loss, establishing a link between an observed CDS term structure and the risk-neutral PD term structure. In our framework, the risk-neutral PD term structure is characterized by a Weibull survival function, which we estimate from the observed CDS term structures of each issuer. For more details, see Section 3.3.

¹² This concept is used widely in pricing many different types of options; it is not limited to securities with credit risk. Please see Gisiger (2009) for an overview of the concept.

¹³ Perhaps the term risk neutral PDs is a misnomer, as it seems to suggest that the term doesn’t account for risk. Oldrich Vasicek used the term “risk-adjusted probability of default,” which could have been more intuitive.

Component 2

The second key model component involves the translation between risk-neutral default probabilities and physical probabilities using market price of risk parameters. This translation involves a simple formula derived from two theoretical constructs. The first is a structural definition of default, under which a firm defaults when its market value falls below its debt obligations. The second is the Capital Asset Pricing Model, which describes the basic trade-off relationship between risk and return.

Conventional tools used for translating CDS spreads into default probabilities typically yield a risk-neutral default probability, unless there is an adjustment for risk aversion. Physical PDs are needed to estimate the distribution of future losses of a credit portfolio and to determine required economic capital. Since risk-neutral PDs are generally much higher than physical PDs, they are not a suitable replacement for physical PDs. For example, Figure 7 shows that on June 30, 2009, for 1,214 North American corporations with CDS spreads, the average ratio between risk-neutral and physical PDs was 2.67.

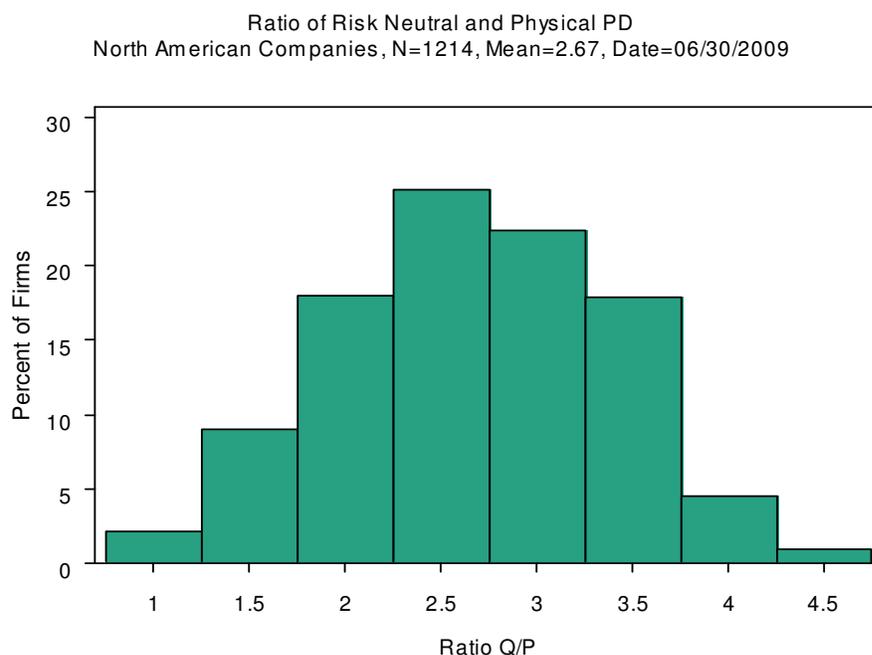


Figure 7 Ratio of Risk-neutral and Physical PD

In the remainder of this section, we describe these two model components and the estimation of the model. We also discuss how the framework can be used to produce a CDS-implied EDF credit measure from a CDS spread and a FVS from an EDF credit measure.

3.3 Deriving Risk-neutral Default Probability

Based on an LGD assumption, a term structure of CDS spreads can be related to a term structure of risk-neutral default probabilities by equating present values of the default and fee legs of CDS contracts under the risk-neutral assumption. Since a term structure of CDS is not sufficient to fully specify the full term structure of risk-neutral PDs, we make the assumption that the risk-neutral survival function is Weibull, which allows us to express spreads in the form:

$$s(t) = \phi(t, h_0, h_1; \delta_t) LGD \quad (1)$$

Where $s(t)$ represents the spread of an t -year CDS contract, LGD is the expected loss rate if default happens, δ_t is the default-free discount curve, h_0 and h_1 are Weibull parameters characterizing risk-neutral default probability term structure so that t -year risk-neutral PD is

$$Q_t = 1 - \exp\left(-(h_0 t)^{h_1}\right) \quad (2)$$

We can see that CDS spread is linear in LGD; the same as in the conventional approximation where spread equals PD multiplied by LGD. This equation is much richer than the “spread equals PD x LGD” approximation in that the full term structure of risk-neutral default probability is accounted for, as are all contingent future cash flows.

We use these formulas for two purposes. First, we estimate the two Weibull parameters most consistent with the CDS term structure and a particular LGD. These two parameters yield a full term structure of risk-neutral default probabilities. Second, we compute a CDS spread term structure from a term structure of risk-neutral PDs and an LGD value.

3.4 Bridging Risk-neutral and Physical Default Probability

The difference between risk-neutral PD and physical PD is driven by the risk premium. The risk premium is determined by the market price of risk, the level of systematic risk, as well as the tenor of the contract. Motivated by Black-Scholes/Merton's structural framework (1974), we translate between a physical and a risk-neutral default probability using the following formula:

$$Q_t = N\left(N^{-1}(P_t) + \lambda \rho \sqrt{t}\right) \quad (3)$$

Where

t :	time horizon
Q_t and P_t :	cumulative risk-neutral and physical default probabilities
λ :	the market price of risk, or the Sharpe ratio
ρ :	the correlation between the asset return of the issuer and of the market
N and N^{-1} :	the cumulative Normal distribution function and its inverse function

While the derivation of the formula involves a number of assumptions, we do not find that transformations based on alternative assumptions were significantly different, once calibrated to the data.

3.5 CDS-implied EDF Credit Measures and Fair-value Spreads

The CDS-implied EDF framework provides a bridge between credit spreads and EDF credit measures. In this section we present an approximation of this translation that is similar to the actual translation, but considerably simpler. This approximation is most accurate when spreads and default probabilities are low.

If the market price of risk and the Sector LGD are known, the approximation for the spread implied by the EDF (what we term the *Fair-value Spread*) is given by the following two formulas:

$$Q_5 = N\left(N^{-1}(P_5) + \lambda\rho\sqrt{5}\right) \tag{4}$$

$$\text{Spread} \approx \text{LGD} \times Q_5 / 5$$

where P_5 and Q_5 are the physical and risk-neutral five year cumulative default probabilities, respectively. Note that we use the cumulative year EDF for P_5 .¹⁴

If the spread is known, we can approximate the CDS-implied EDF using these two formulas:

$$Q_5 \approx 5 \times \text{Spread} / \text{LGD} \tag{5}$$

$$P_5 = N\left(N^{-1}(Q_5) - \lambda\rho\sqrt{5}\right)$$

In the actual implementation, we do not use the approximation: $\text{Spread} \approx \text{LGD} \times Q_5 / 5$, but we work with the full term structure of risk-neutral probabilities. In computing the fair-value spreads, we estimate the term structure of spreads based on the population of firms with comparable risk levels.

3.6 Model Estimation

The key intermediate parameters of the CDS-implied EDF model are the market Sharpe ratios and the sector LGDs. We estimate market Sharpe ratios by different geographic regions and by different asset classes. This feature is motivated by observing, that for companies with similar EDF levels, we find clusters of CDS spreads by region and by asset classes. Across regions, controlling for EDF levels, Japanese companies tend to have much lower CDS spreads; across asset classes, controlling for EDF levels, investment grade companies tend to have lower spreads than high yield entities.

We estimate LGDs by industrial sectors for each region using a two-step process. First, with an initial LGD assumption of 60%, we estimate the market Sharpe ratio that brings overall consistency between EDF levels and spread levels for the region and for the asset classes. Spreads in certain sectors can still be systematically higher or lower than spreads predicted from EDF credit measures using the calibrated market Sharpe ratios and the 60% LGD assumption. We then remove these systematic differences by calibrating sector LGDs, so that within each region and sector, spreads and FVSs are, on average, consistent.

Figure 8 shows how risk premiums differ by region for investment grade firms. In the time series, we can see that risk premiums increased significantly during the “great recession” as retail investors hoarded cash and capital markets around the globe experienced a severe credit crunch. These are signs that investors are more risk averse and ask for higher compensation per unit of risk. In the cross-section, we observe a big difference in the market price of risk between the Japanese market and the rest of the world.

¹⁴ We bridge the EDF and the QEDF at the five year horizon as the CDS market is the most liquid at the five year horizon. For ease of exposition, we are using the approximation that the annualized Q is one-fifth the five year cumulative Q , a good approximation for small default probabilities.

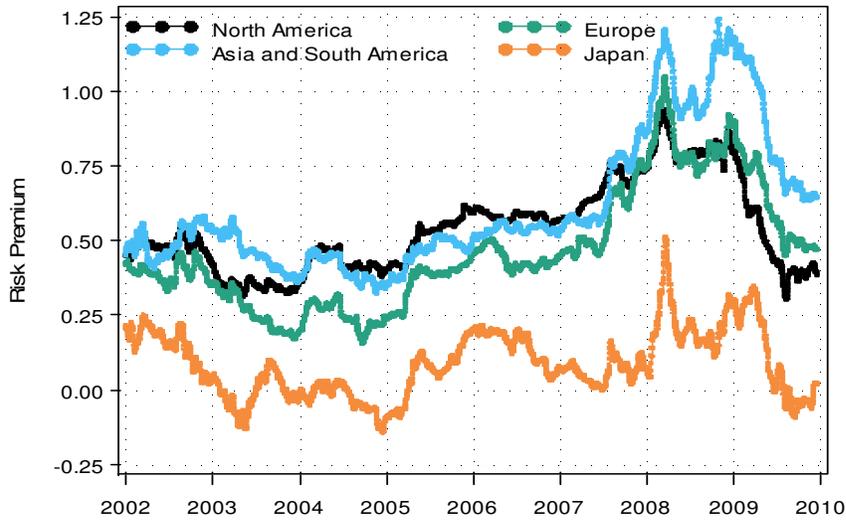


Figure 8 Risk Premium By Region Category for Investment Grade Firms

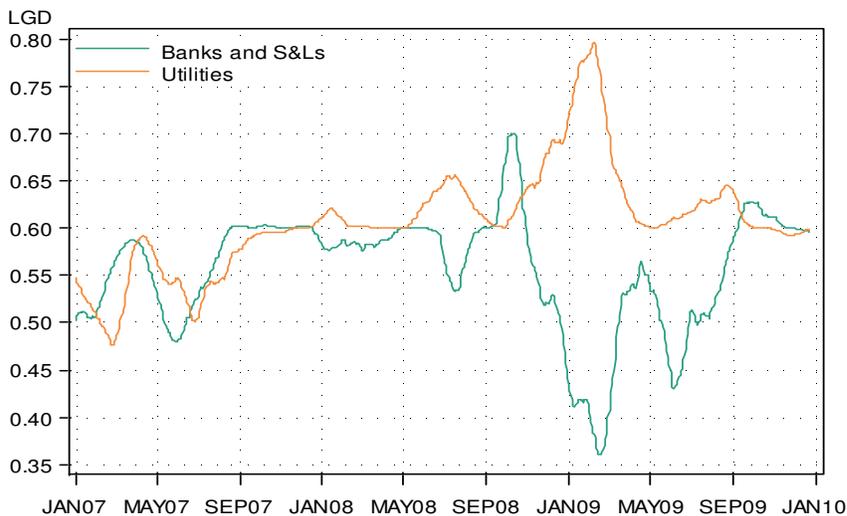


Figure 9 North American Sector LGDs

Figure 9 illustrates LGD for two sectors: North American Banks and S&Ls and North American Utilities. During the beginning of 2009, the high sector LGD for North American Utilities reflects elevated spreads for North American Utilities relative to spreads in other sectors and similar EDF credit measures. A rapid increase in LGD typically reflects spreads increasing in the sector without a comparable increase in the EDF, or EDF credit measures decreasing in the sector without a comparable decrease in spreads.

Within each region, sector LGDs are centered around 60%. When the underlying LGDs systematically deviate from the assumed initial value of 60%, the difference is compensated for by estimated Sharpe ratios. Thus, the estimated LGD for a sector mainly captures the sector's LGD relative to other sectors, not the time-series variation.¹⁵ An LGD of 60% is often assumed by market participants for corporate issuers, as we have observed from data contributed by dealers to CDS data vendor Markit.

CDS-implied EDF credit measures for each individual issuer are computed using its CDS spreads and estimated model parameters specific to the issuer's geographical location, rating status, and business sector. The exception is the sovereign CDS market, where an LGD of 75%, as opposed to 60%, and the North American Sharpe ratios, as opposed to regional Sharpe ratios, lead to more reasonable CDS-implied EDF credit measures. We elaborate further on this application in Section 2.

3.7 Model Validation

We validate the model from various angles. First, we validate the estimation of the model. In order to check that our implementation is working as intended, we verify that the geometric average of the FVS is consistent with the geometric average of the actual spread over time, and that the geometric average of the CDS-implied EDF is consistent with the geometric average of the actual EDF over time. We next look at the cross-sectional correlations of both the FVS and spread and the CDS-implied EDF and the EDF over time for different regions. As a third test on the model, we look at the model's implications for the sensitivity of credit returns to equity returns and compare them with empirically measured counterparts. We also examine the behavior of differences between the fair-value spread and the actual spread over time. We have already reported results on cumulative accuracy profiles that result from using the framework in Section 2.2.

¹⁵ Sector LGDs are not a pure measure of relative LGD expectations. Spread levels in a sector reflect default risks and expected LGDs as well as factors not included in the model, such as liquidity levels, hedging demand, and ease of diversifying credit risk. Sector LGDs capture these factors and move the average implied CDS spreads toward the observed average spread within each sector. Sector LGDs also help remove systematic differences between the market's view of credit risk in a sector and the EDF credit measure.

3.7.1 Comparison of Geometric Means

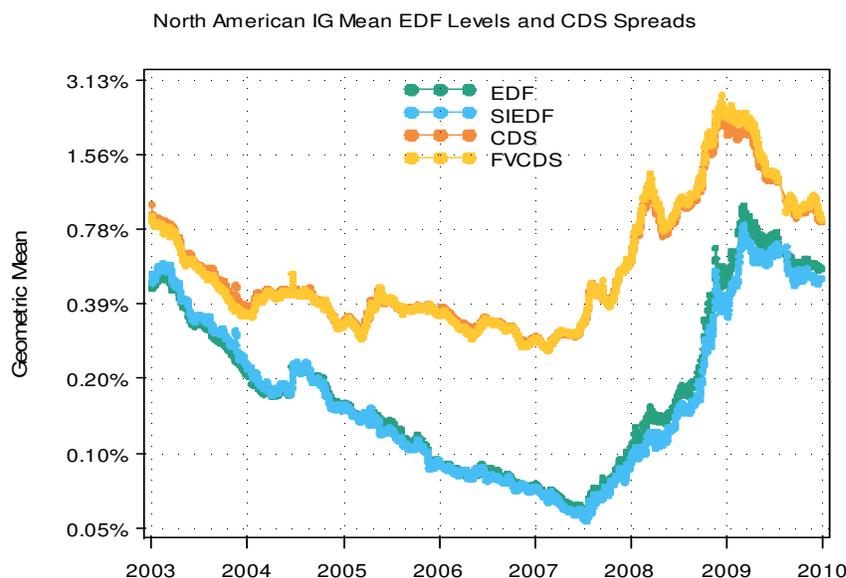


Figure 10 North American Investment Grade Mean EDF Levels and CDS Spreads

As expected, our model estimation process produces CDS-implied EDF credit measures comparable to EDF credit measures, and produces fair-value spreads comparable to observed CDS spreads, when averaged across a regional sample of given asset class (i.e., Investment Grade or High Yield).¹⁶

Figure 10 demonstrates such behavior with the sample of North American Investment Grade companies. The geometric averages of CDS spreads and fair-value spreads are very close on the common sample at any given time; the geometric averages of EDF levels and CDS-implied EDF levels are also very close on the common sample at any given time. One reason that the matches are not exact is that the sample is slightly different from the calibration sample. Since the averages are very close for each cut of the data on the common sample, fair-value spreads are comparable to CDS spreads. Similarly, the interpretation of a CDS-implied EDF level is comparable to that of an EDF level.

3.7.2 Cross-sectional Correlation

The second aspect of the validation is to examine the cross-sectional correlations between CDS spreads and EDF implied FVS. As shown in Figure 11, the correlation has varied around a high level of 60% for North American Investment Grade companies. We notice that such correlations tend to be lower in benign times, probably because there is a larger, non-credit component to spreads. For European Investment Grade companies, we see the correlations are comparable to their North America counterparts until the most recent recession. Note that the correlation between CDS and FVSs in Europe has recently increased. This finding is the second aspect of the validation.

¹⁶ We calibrate the model using logarithm of all PD variables and spread variables. Therefore by “average,” we refer to geometric average.

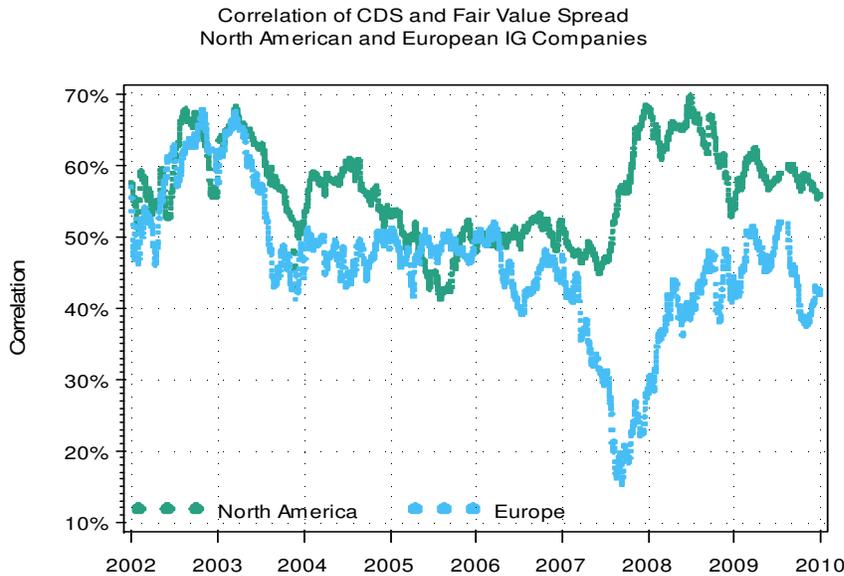


Figure 11 Correlation of CDS and FVS Through Time

3.7.3 Sensitivity of Credit Returns to Equity Returns

The third aspect of validation focuses on the prediction of the sensitivity of debt prices to equity prices.¹⁷ Our measure of this sensitivity is the debt to equity hedge ratio (i.e., the percentage change in debt prices per percent change in equity prices). The hedge ratios can be useful to assess value changes in loan or bond portfolios when such credit instruments are not actively traded. When we compare the hedge ratio predicted from CDS-implied EDF framework to the regression based hedge ratio, we find that they are reasonably consistent, as seen in Figure 12. In the figure, we first sample public companies that have both equity and CDS into different EDF groups. Next, we compute average hedge ratios implied by our theoretical framework; third, for each group we regress bond returns on equity returns to obtain the empirical hedge ratio. Finally, we compare EDF implied hedge ratios to empirical hedge ratio.

In Figure 12, hedge ratios based on our structural framework are compared to hedge ratios derived from regression on data for North American companies bucketed by EDF5. We can see that for safe or low EDF companies, both modeled hedge ratio and empirical hedge ratio are lower, and for risky or high EDF companies, both hedge ratios are much higher. This finding is intuitive because for safe companies, a modest decline in asset value has little impact on credit risk and, hence, little impact on bond returns. For risky companies, on the contrary, a modest decline in asset value increases credit risk and both equity investor and bond investors will experience negative returns.

¹⁷ Schaefer and Strebulaev (2008) note that structural models provide accurate descriptions of the sensitivity of bond returns to equity returns.

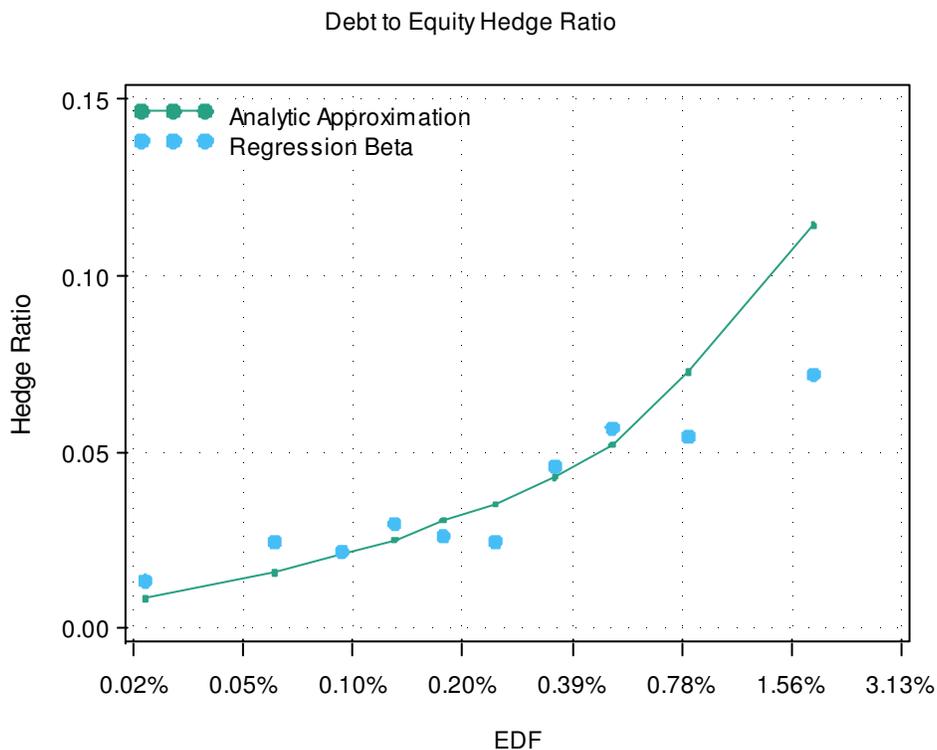


Figure 12 The Model-implied Sensitivity of CDS Returns to Equity Returns Compared to Empirical Sensitivity

3.7.4 Behavior of Differences over Time

The fourth aspect of the validation examines whether or not differences between the FVS and the actual spread persist over time. The Fair-value Spreads produced by the model can be higher than observed CDS spreads for some firms and lower for other firms. When FVS and CDS spreads are different, it is possible that either the CDS spread is too high (or too low), or its equity price is too low (or too high). Our empirical analysis shows that the differences between CDS and FVS have a large transitory component (i.e., they tend to dissipate over time). Such a reduction in differences indicates convergence between the CDS spread and equity valuations. Such pricing behavior may be indicative of a potential trading strategy for an asset manager.¹⁸

Specifics of our empirical analysis follow. Each month m , we rank companies into five groups according to their differences between CDS and FVS. For each group, we track average price differences so that

$$L_{m,m+t,k} = \frac{1}{I_{m,k}} \sum_{i=1}^{I_{m,k}} \log(CDS_{i,m+t} / FVS_{i,m+t}) \quad (6)$$

Where $I_{m,k}$ is the number of firms in group k , $1 \leq k \leq 5$ and $0 \leq t \leq 36$. For example, $L_{2003/12,6,5}$ (with m ="December 2003", $t=6$ and $k=5$) represents June 2004 average pricing difference of the group of companies that had the highest pricing differences in December 2003. Figure 13 illustrates the average log (CDS/FVS) for five company groups formed at the end of 2003 (i.e., m is December 2003).

¹⁸ Li and Zhang, 2010.

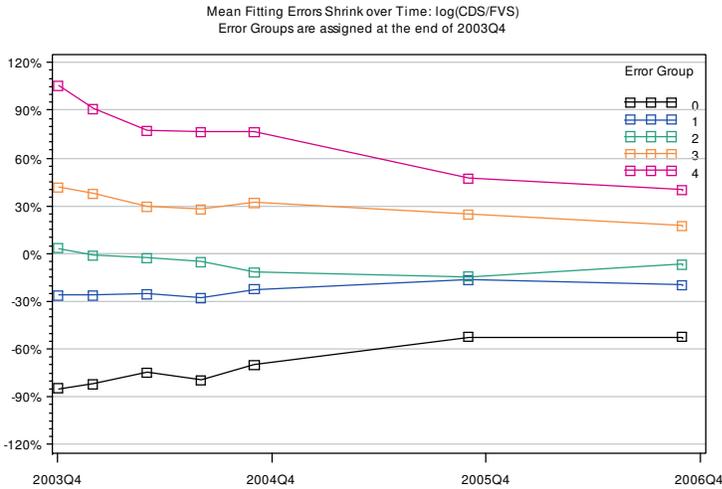


Figure 13 Convergence in Differences between CDS and Fair-value Spreads for Portfolios Formed at the End of 2003

We repeat this exercise of difference tracking for all months m and compute the average $L_{m,m+t,k}$ across all months m for each group k and tracking horizon t :

$$\bar{L}_{t,k} = \frac{1}{M} \sum_{i=m}^M L_{m,m+t,k} \quad (7)$$

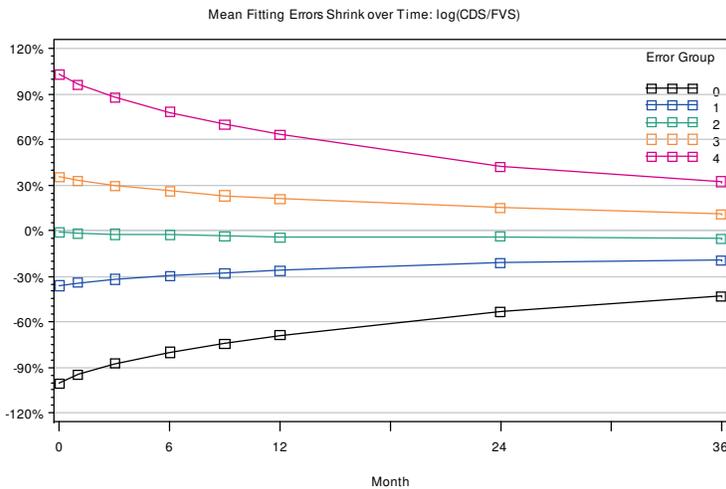


Figure 14 Convergence in Average Differences between CDS Spread and Fair-value Spread

As seen from Figure 14, $\bar{L}_{t,1}$, the average differences for firms initially having higher CDS than FVS declines with time t , indicating that if a company initially had CDS too high relative to FVS, either its CDS spread would decline or its equity price would decline (negative equity returns would lead to higher EDF and higher FVS). This finding indicates that a hedging portfolio that longs CDS and shorts equity could potentially be profitable on average. The

opposite is true for companies that had initial CDS too low relative to FVS; a hedging portfolio that shorts CDS and longs equity would also be potentially profitable.

Further analysis shows that the reduction of pricing differences is caused by movements of both CDS spreads and equity prices. In Figure 15, we report average CDS and FVS for companies in two portfolios with extreme initial pricing differences. The left panel shows companies that initially had FVS much higher than CDS spreads. For these companies, their CDS consistently increase, their FVS consistently decrease, and their difference decreases from the initial 260 basis points to about 100 basis points three years later. The right panel demonstrates a similar story, in the opposite direction, for companies that initially had FVS much lower than CDS spreads—FVS consistently increase through out the tracking period of three years, and CDS spreads decrease consistently for first two years.

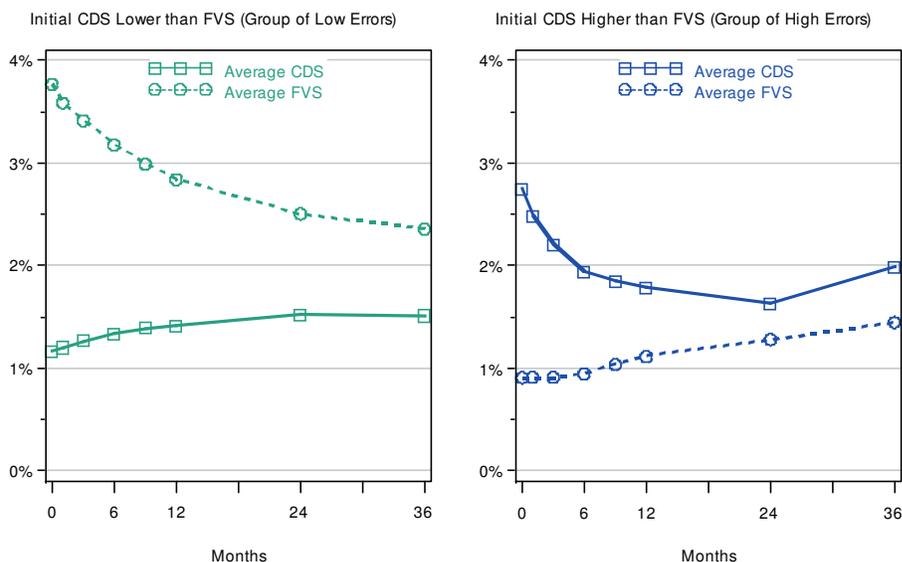


Figure 15 Interaction Between CDS Spread and FVS Leads to Their Convergence

In summary, this framework provides CDS-implied EDF credit measures on the same average level as EDF credit measures, and FVSs on the same average level as CDS spreads, bringing information from the CDS market and the equity market to a level playing ground with a reasonably high cross-sectional correlation. The model predicts hedge ratios consistent with empirical values, thus allowing assessment of changes in credit portfolio values using equity value changes, a useful application for illiquid credit exposures.

The model also identifies differences in opinion between the CDS market and the equity market, and as the differences tend to converge, this allows asset managers to explore profitable trading strategies from such differences. Additionally, multiple measures of default risk are shown to improve default prediction, as reported in Section 2.2.

4 Applying the Framework to Sovereigns

Like corporate debt, sovereign debt also carries default risk. Even countries with the most developed economies carry risk, as illustrated by the five-year CDS spread of the U.S. government and UK government in Figure 16. The CDS market for sovereigns has been developing, along with the corporate CDS market. The number of sovereign countries covered by the CDS market has been steadily increasing, from the low 30s in 2001 to more than 80 in

recent years, as shown in Figure 17. In this section, we describe how we adapt the framework for extracting the physical default probability of a corporate from the credit spread of the corporate to the sovereign.

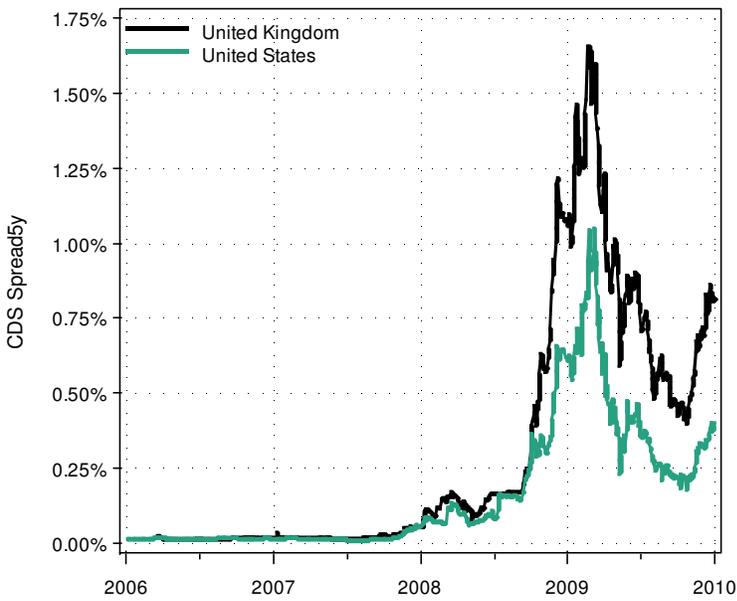


Figure 16 CDS Spreads for UK and U.S. Sovereigns

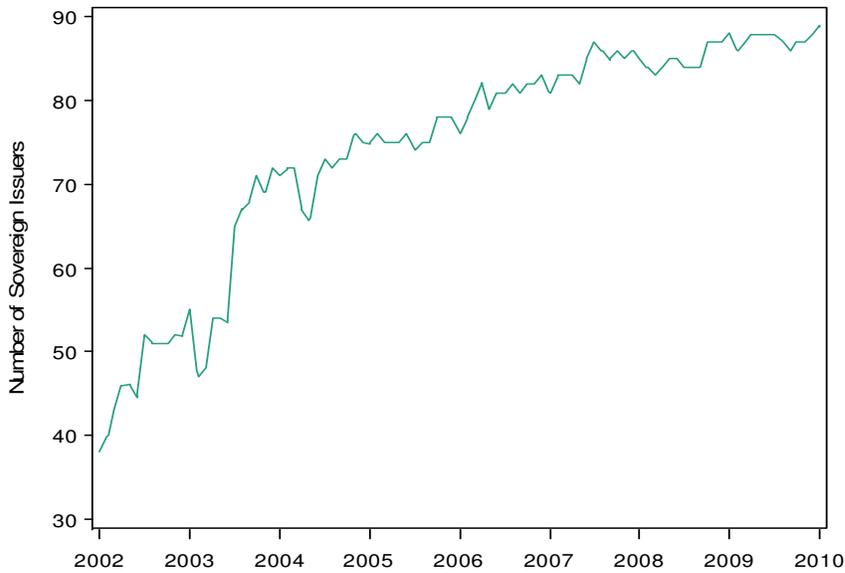


Figure 17 Number of CDS Quotes on Sovereigns Over Time

4.1 Implementing the Framework

To illustrate, we repeat the following simple approximation of the framework to convert a spread into a physical default probability:

$$Q_5 \approx 5 \times \text{Spread} / \text{LGD} \quad (8)$$

$$P_5 = N\left(N^{-1}(Q_5) - \lambda \rho \sqrt{5}\right) \quad (9)$$

Where P_5 and Q_5 are the cumulative five year physical and risk-neutral default probabilities, respectively. In principle, the same framework that converts a spread into a CDS-implied EDF for a corporate can be applied to sovereigns. Relative to corporates, however, we make different choices for sector LGD and market Sharpe ratio. In the case of corporates, we estimate the market price of risk using the sample of corporates with both an EDF and a CDS spread within an asset class (investment grade versus speculative grade) and a region. We also estimate the sector LGD using a sample of corporates in the sector/region with both EDF and a CDS spread. For sovereigns, such a sample does not exist, as sovereigns do not have EDF credit measures, since they do not have publicly traded equity.

We choose to apply the market Sharpe ratio estimated from the North American corporate sample across all sovereign issuers, which implies that an investment grade sovereign issuer would use the same market Sharpe ratio as a U.S. investment grade corporate, and a high yield sovereign issuer would use the same market Sharpe ratio as a U.S. high yield corporate. Using different market Sharpe ratios for different rating classes is motivated by market segmentation; using North American Sharpe ratios instead of regional Sharpe ratios is motivated by the observation that investors of sovereign foreign currency debt most likely come from developed markets, and the North American debt market is the most liquid in terms of trading depth.

The recovery process of a sovereign default differs from corporate default and is more uncertain. When a sovereign issuer fails to honor its debt obligations, no court enforces liquidation. Often, existing debt obligations are restructured with new debt, sometimes with lengthy delays and significant losses to investors. When pricing a CDS contract for a corporate, the market convention is that the baseline LGD is 60%. For sovereigns, it is often 75%. These figures are clearly seen from the dealer quoted LGDs, shown in Figure 18. Therefore, we assume an LGD of 75% for all sovereign issuers.

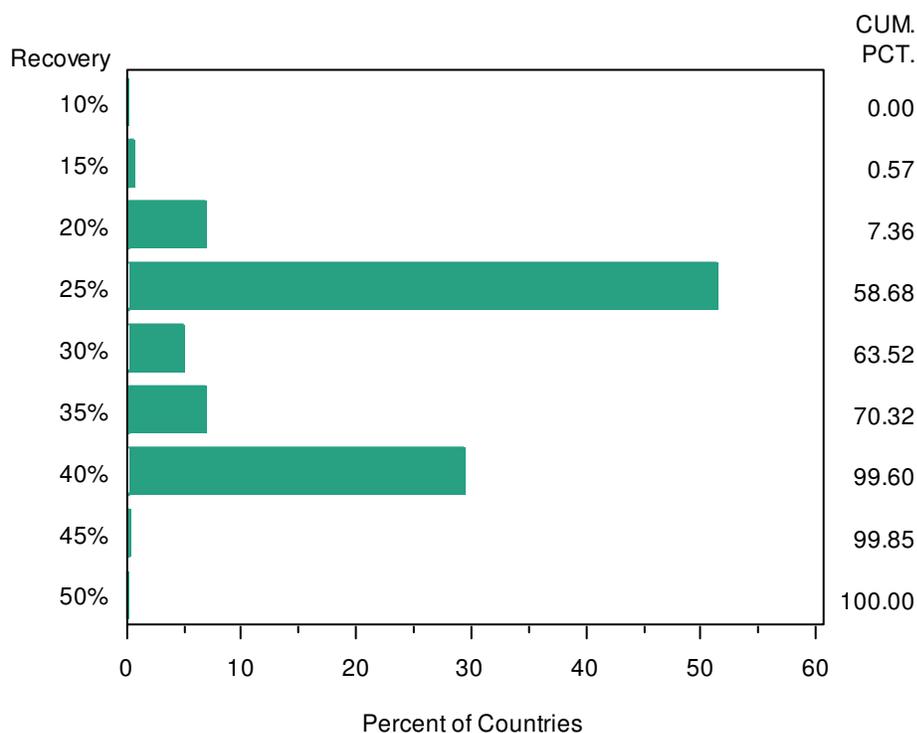


Figure 18 Sovereign Debt Recovery Distribution

4.2 Validation

Validating the framework for sovereigns is challenging because we do not have equity-based EDF credit measures to benchmark against, nor do we have a large number of default events with which to validate. Therefore, although the framework of the CDS-implied EDF model for sovereigns is similar to that of corporates, our validation approach is different.

We validate our implementation of the framework on sovereigns in three ways.

1. Compare sovereign CDS-implied EDF credit measures to the corresponding rating implied EDF to ensure that their levels are comparable, and that they are well correlated.
2. Verify that sovereign CDS-implied EDF credit measures are low relative to the CDS-implied EDF for corporates in the same region, as expected.¹⁹
3. Examine the time series of CDS-implied EDF credit measures for the three sovereigns that have defaulted with liquid CDS spreads to ensure CDS-implied EDF credit measures provided a timely early warning relative to ratings.

¹⁹ For example, Moody's Investors Service uses the concept of a "Country Ceiling," which, under certain circumstances, caps the ratings of issuers within a country. For more information see "A Guide to Moody's Sovereign Ratings," Moody's Investors Service, December 2008.

4.2.1 CDS-implied EDF Credit Measures Compared to Rating Implied EDF Credit Measures

We first establish the suitability of rating-implied EDF credit measures as a benchmark for sovereign CDS-implied EDF credit measures. We then use rating-implied EDF credit measures to benchmark both level and rank ordering for sovereign CDS-implied EDF credit measures. We show that our use of 75% LGD produces CDS-implied EDF levels comparable to Rating Implied EDF levels. Moreover, our consistent use of North American market Sharpe ratios results in a high correlation between CDS-implied EDF credit measures and Rating Implied EDF credit measures.

Rating Implied EDF credit measures are computed as the median EDF for corporates of a given rating, as produced in Moody's Analytics CreditEdge. Thus, the suitability of Rating Implied EDF credit measures as a benchmark for sovereign CDS-implied EDF credit measures depends on the consistency of ratings between corporates and sovereigns. Evidence for this consistency is given by the similarity of CDS spreads for corporates and sovereigns of the same rating, as shown in Figure 19, which compares median CDS spreads for corporations and for sovereign issuers by rating groups. As of the end of 2009, the sovereign spread was higher than that of the corporate spread for A, and Aa, Baa credits, but lower for Ba and very similar for Aaa credits. The difference in spreads between corporates and sovereigns of the same rating category is much smaller than the difference between spreads for an Aa versus a Ba credit for either corporates or sovereigns.

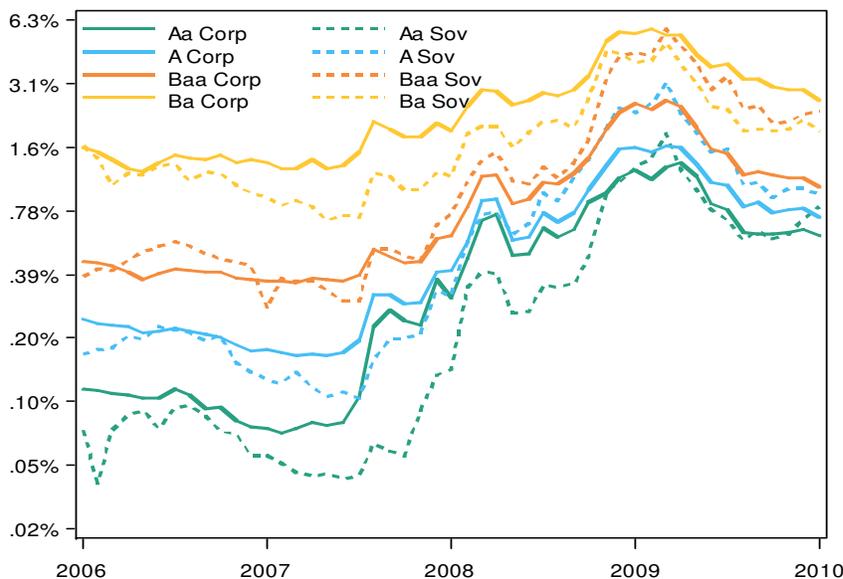


Figure 19 Median CDS Spread by Rating

Figure 20 compares the level of the CDS-implied EDF for sovereigns with their rating-implied EDF credit measures for different LGD choices. It shows that if we use the same 60% average LGD for corporations and sovereigns, CDS-implied EDF credit measures for sovereigns would be too high when measured by either the geometric mean, the mean, or the median relative to the Rating Implied EDF. The assumption of a 75% LGD lowers the CDS-implied EDF credit measures of sovereigns and makes them more comparable to the Rating Implied EDF credit measures.

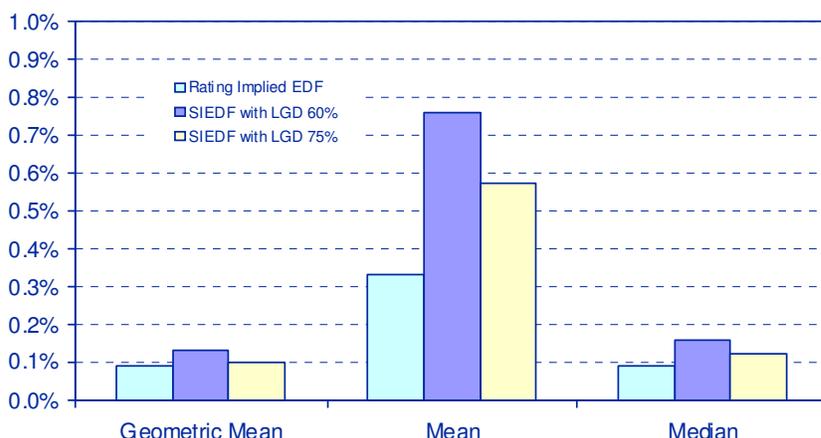


Figure 20 Comparing Average CDS-implied EDF to Average Rating-implied EDF

Overall, the correlation between sovereign CDS-implied EDF credit measures and rating-implied EDF credit measures is about 85%, compared with 72% of a random sample of corporations with the same sample size.

One way to test whether the level of a sovereign CDS-implied EDF is reasonable is to see if it resides among the lowest CDS-implied EDF credit measures of corporate entities within the same country. This is because we expect the sovereign issuer to be among the safest issuers from the country. When a sovereign issuer experiences credit difficulties, the entire national economy may also experience systematic problems. Such problems usually affect corporations as well. In addition, there have been cases where a sovereign government was having trouble meeting its foreign debt obligation and issued administrative orders to limit cross-border financial payments, prohibiting corporations from making contractual payments, even when they were economically capable of doing so. Therefore, we expect a PD measure for a sovereign to be comparable to, if not better than, those of the best corporations in the country.

Prior to the credit crisis in 2007, for all countries with CDS-implied EDF credit measures for at least 20 corporates except Sweden and Switzerland, the sovereign CDS-implied EDF was lower than the fifth percentile of the corporate CDS-implied EDF credit measures within the country more than 80% of the time. During the credit crisis, sovereign CDS-implied EDF credit measures were more elevated.

Figure 21 shows the case of CDS-implied EDF for the government of Chile and Chilean corporations. With the assumption of a 75% LGD, the CDS-implied EDF of the government of Chile is effectively reduced, making it the least likely to default among all Chilean issuers.

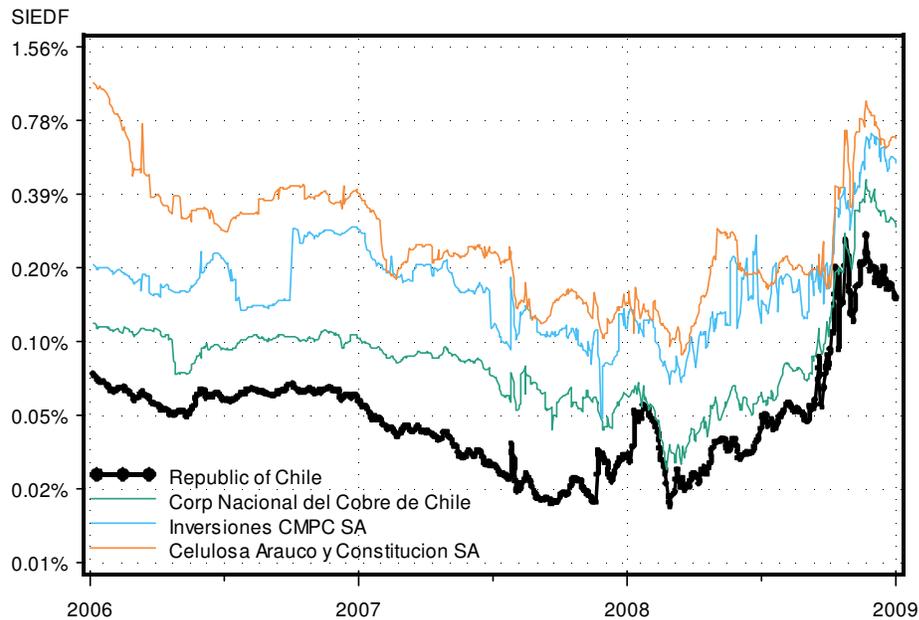


Figure 21 CDS Spread of Chilean Corporations and the Sovereign

4.2.2 CDS-implied EDF Credit Measures for Actual Defaults

To date, there have been three sovereign defaults for which there has been a liquid CDS market prior to default: Ecuador, Uruguay, and the Dominican Republic. In all three cases, both the CDS spread and the CDS-implied EDF were elevated prior to default. Relative to rating, the CDS-implied EDF of all three provided a more dynamic indicator of risk.

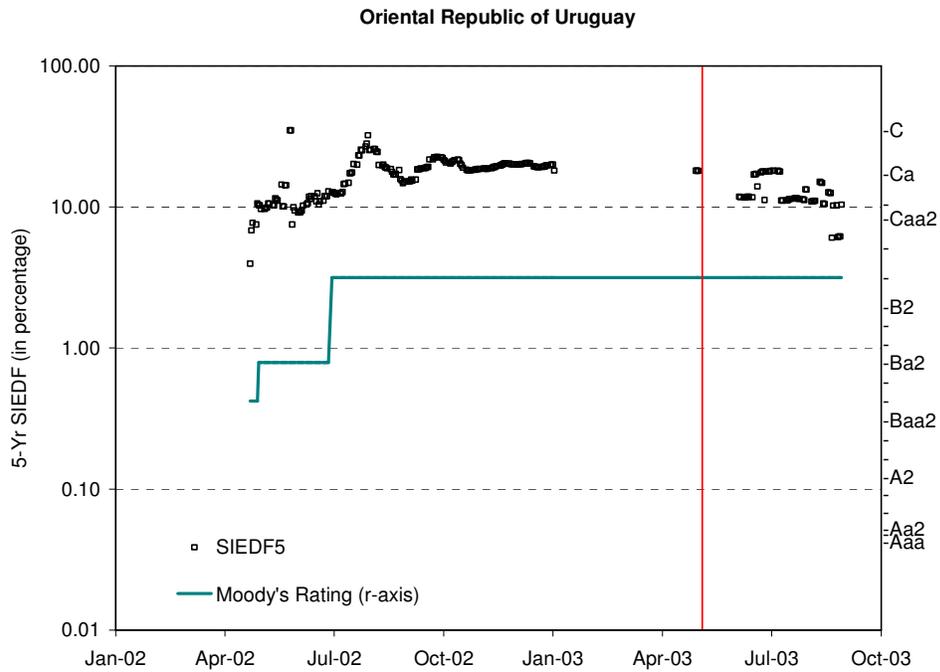


Figure 22 Uruguay's CDS-implied EDF and Rating, Pre- and Post-default

In May 2003, while facing a currency crisis, Uruguay completed a distressed exchange by extending the maturities of existing bonds. Moody's Investors Service interpreted this event as a contagion default from the Argentinean default of 2001.²⁰ The CDS-implied EDF was inflated at the time of default relative to the rating, as shown in Figure 22.

²⁰ See "Sovereign Default and Recovery Rates, 1983-2007," Moody's Investors Service, March 2008.

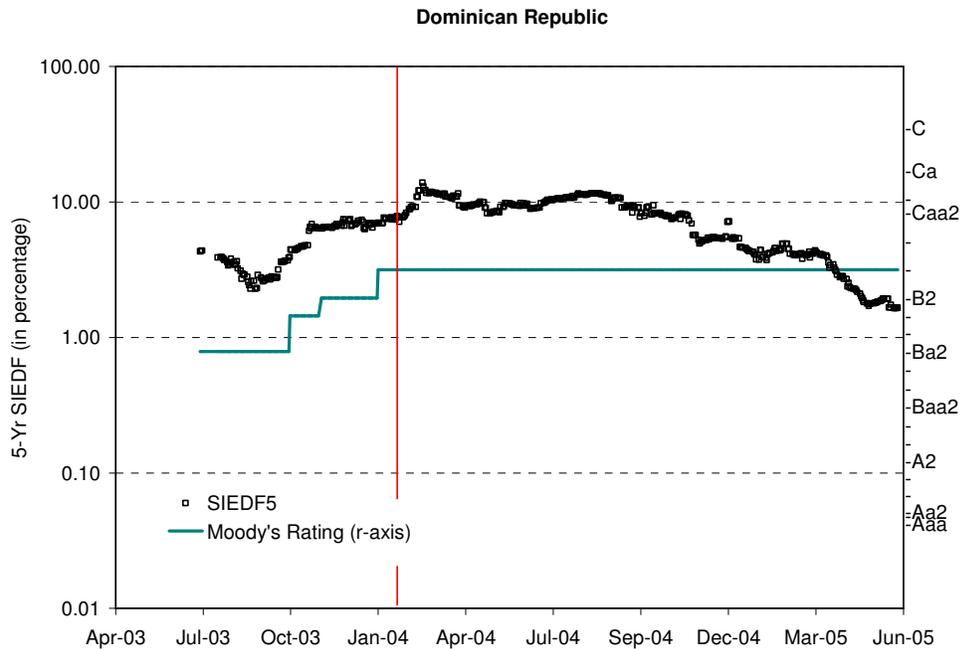


Figure 23 Dominican Republic CDS-implied EDF and Rating, Pre- and Post-default

In January of 2004, the Dominican Republic missed a bond payment but was able to cure the default within a 30-day grace period. After missing a number of payments during the following year, in April of 2005, they completed a distressed exchange. The CDS-implied EDF was very inflated throughout this process, as shown in Figure 23.

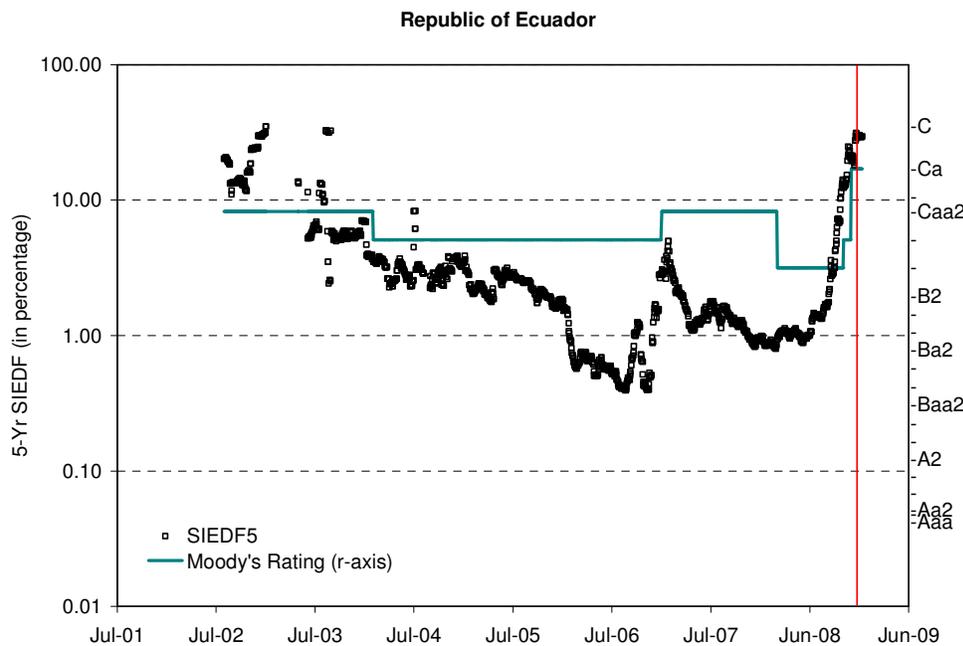


Figure 24 Ecuador's CDS-implied EDF Prior to Default

In December of 2008, the Ecuadorian government announced that two of its sovereign debt securities were “illegal.” The government formulated a restructuring plan with a severe haircut, shown in Figure 24. Moody’s Investors Service describes this default as an issue of “willingness to pay” rather than “ability to pay.”²¹ The CDS-implied EDF rose markedly prior to default from relatively low levels earlier in the year.

In summary, sovereign CDS-implied EDF measures provide market-based physical PD measures for sovereigns that can be used together with equity-based EDF measures for corporations. These CDS-implied EDF measures are consistent with credit ratings as indicated by their high correlation with rating implied EDF credit measures. The CDS-implied EDF of a sovereign is typically lower than the CDS-implied EDF credit measures for the safest firms in the country. Further, the CDS-implied EDF of the three cases of sovereign defaults have been more dynamic and the rating. Applying the CDS-implied EDF to sovereigns allows the risk manager to compare the default risk of a sovereign with its corporate counterparts on a comparable basis.

5 Summary

In this paper, we present a model that links two commonly used risk metrics: physical default probabilities and credit spreads. With this model, we perform credit assessment using objective and forward-looking information from both the equity and CDS markets. The model provides CDS-implied EDF credit measures that can be compared directly with EDF credit measures; the model also provides Fair-value Spreads that can be compared directly with observed CDS spreads.

We construct the link between a physical PD and a spread using an intermediate risk metric, the risk-neutral PD. Going from a Spread to an EDF is accomplished in two steps. In the first step, the term structure of CDS spreads is translated to a risk-neutral PD term structure that balances the expected value of the premium payments against the expected value of the recovery payments of CDS contracts, given a calibrated LGD parameter. In the second step, the physical PD is derived from the risk-neutral PD using a calibrated risk premium parameter. Both the LGD parameters and the risk premium parameters are calibrated empirically on an ongoing basis. We calibrate the LGD parameters by geographical region and industry sectors, and we calibrate the risk premium parameters by geographical region and rating types (i.e., Investment Grade and High Yield). This calibration ensures that, at any given time, average levels of CDS-implied EDF and of EDF are comparable, and average levels of fair-value spread and of observed CDS spreads are comparable.

It is useful to compare equity market information and CDS market information directly, on level playing fields. On the one hand, information from these two markets is complementary. Many companies have an EDF measure but are not covered by the CDS market. In such cases, an EDF implied FVCDS can be used for marking-to-market illiquid credit instruments. On the other hand, in a credit portfolio there may be obligors who are covered by the CDS market but do not have an EDF measure. In this case, EDF credit measures can be used together with EDF credit measures for public firms. In addition, when both CDS information and equity information are available for a company, this information can be used to reinforce credit opinions when EDF and CDS-implied EDF are consistent or to trigger an additional research action when they are inconsistent. We show that one simple way of combining EDF and CDS-implied EDF—to take the more conservative assessment of the two—leads to superior performance in predicting defaults.

In summary, the model framework integrates equity market information and CDS market information. Empirical evidence shows that such integration is valuable in providing useful signals for risk management and credit investment.

²¹ See “Sovereign Default and Recovery Rates, 1983-2008,” Moody’s Investors Service, March 2009.

Appendix A Frequently Asked Questions

This section presents a list of frequently asked questions.

1. Why do we use Weibull survival functions to derive risk-neutral PDs?

In order to solve for the term structure of PDs that set the present discounted value of the default leg of a CDS contract equal to the protection leg, we need to specify a functional form for the hazard function of the firm. The Weibull survival function allows us to control both the level and “slope” of the hazard function. These two degrees of freedom suffice to fit most spread data quite well (as measured by RMSE or variance explained). In addition, the Weibull survival function is analytical tractable, and the function is numerically stable when spreads are added or removed from the term structure.

If the term structure of spreads is either upward-sloping or downward-sloping, we typically see that the Weibull survival function fits the term structures of spreads with a small amount of error, as shown in the top two panels in Figure 25. If the term structure is either “humped” or “bumpy,” then the Weibull term structure will smooth out the “humps” and the “bumps,” and the error will be somewhat larger, as shown in the bottom two panels in Figure 25.

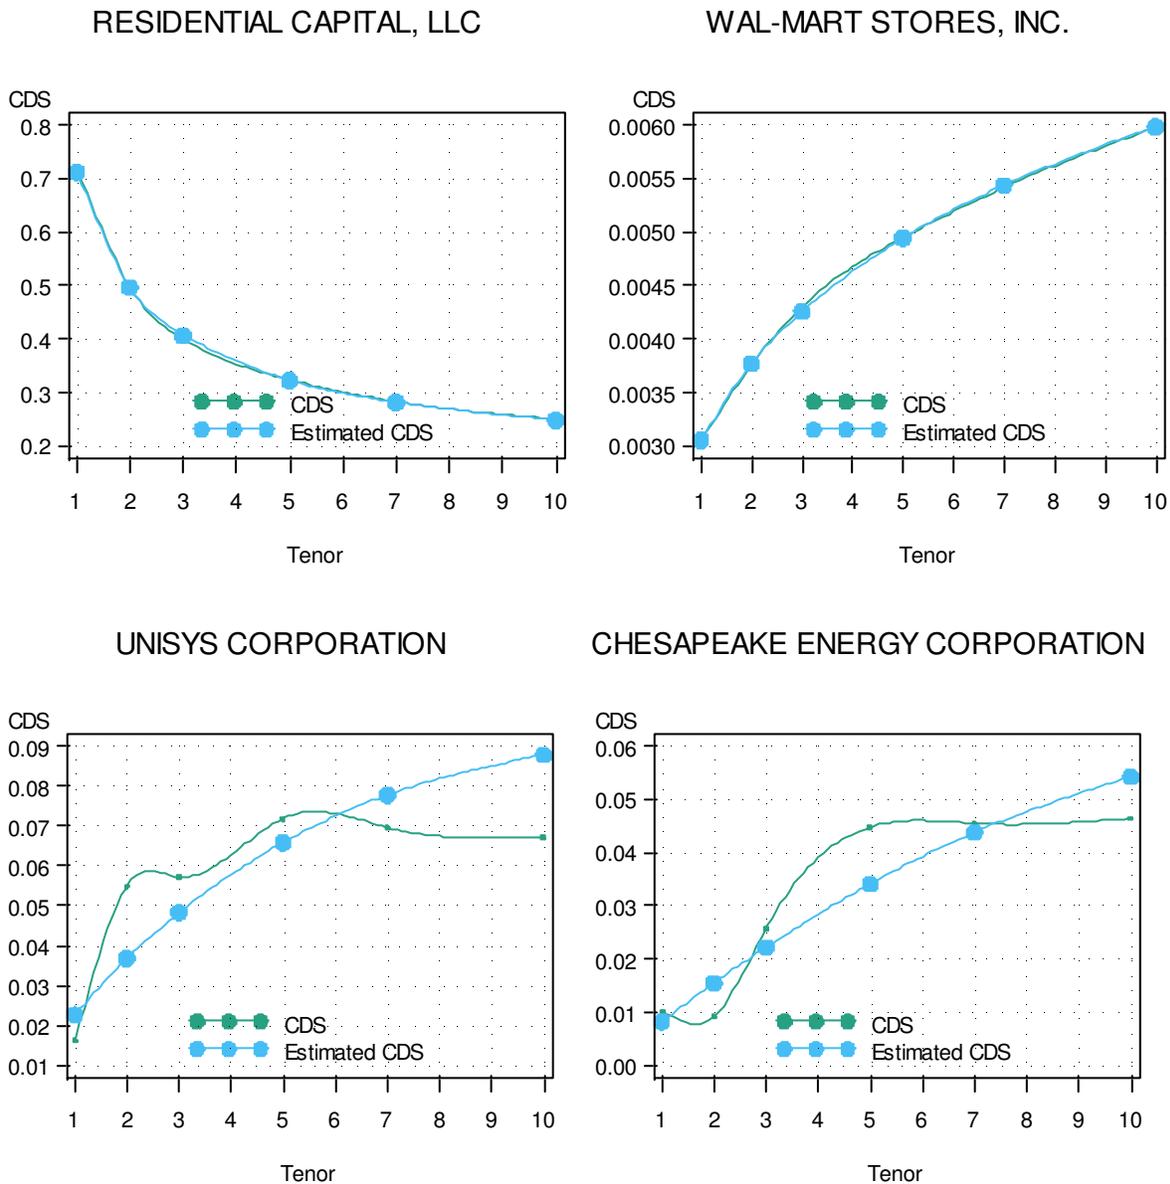


Figure 25 Using a Weibull Survival Function to Fit a Term Structure of CDS Spreads

The Weibull survival function has two parameters that can impact both the “level” and the “slope” of the term structure. We find that most of the variation in spread changes can be explained by the “level component” and “slope component” using a principal component analysis. Principal component analysis is a technique that reduces the dimensionality of the data.

In this context, principal components can be described as follows. First construct a matrix Y , where Y_{ij} is the change in CDS spreads for observation i and tenor j . Each row is an observation and each column is a tenor. An observation is a specific name on a specific day. If the spread at the five-year horizon increases by a large amount, the rest of the spreads are likely to increase by a large amount. Additionally, on any given day, the spread term structure could “steepen” or “flatten.” The so-called first principal component is the linear combination of the columns of Y that explain the largest amount of the variation in matrix Y .

The second principal component is the linear combination of the Y s that explains the largest amount of the residual variation in the Y s and is orthogonal to the first principal component. We find that the first principal component explains 94% of the variation in changes, and the second principal component explains another 4%. The weights used in constructing the first principal component are approximately the same across the different tenors.

Consequently, we interpret the first principal component as representing a “level shift” in the term structure. For the second principal component, the weights for the short tenor spreads have the opposite sign as the weights for the long tenor spreads. Consequently, we interpret the second principal component as representing a “slope shift.” Therefore, a survival function that allows for both a “slope” and a “level” component to the term structure of spreads is capable of capturing most of the variation in spread changes across the term structure.

2. Why do we employ different market prices of risk for different regions and ratings classes?

We use different Sharpe ratios to better match spreads and EDF levels across both regions and ratings classes.

A number of region-specific factors often create discrepancies between regions, where, for a given EDF level, spreads may be much higher or lower on average than in another region. For example, Japan stands out. For a given EDF level, Japanese spreads tend to be considerably lower than those in other regions. Figure 26 shows how Japanese spreads are lower than North American spreads. Historically, the main Japanese banking system followed the practice of banks and parent companies extending credit to companies that would otherwise default. This practice likely leads to lower spreads for the same EDF level. Moreover, the spreads for no restructuring doc clauses (XR), CDS contracts in Japan are particularly low, suggesting that Japanese defaults are more likely to be restructuring events.

Likewise, ratings classes experience variability, despite similar EDF levels. Historically, spreads for High Yield CDS contracts have been higher than Investment Grade CDS contracts, for firms with similar EDF levels. In fact, before 2008, the difference was persistent and substantial. In contrast, the significant systematic differences in spreads for firms of similar EDF levels but different ratings within one of these two broad rating classes are much more limited. One explanation for the differences between these two markets is that they are not fully integrated; many investors in investment grade credit have a limited ability to invest in speculative grade credits.

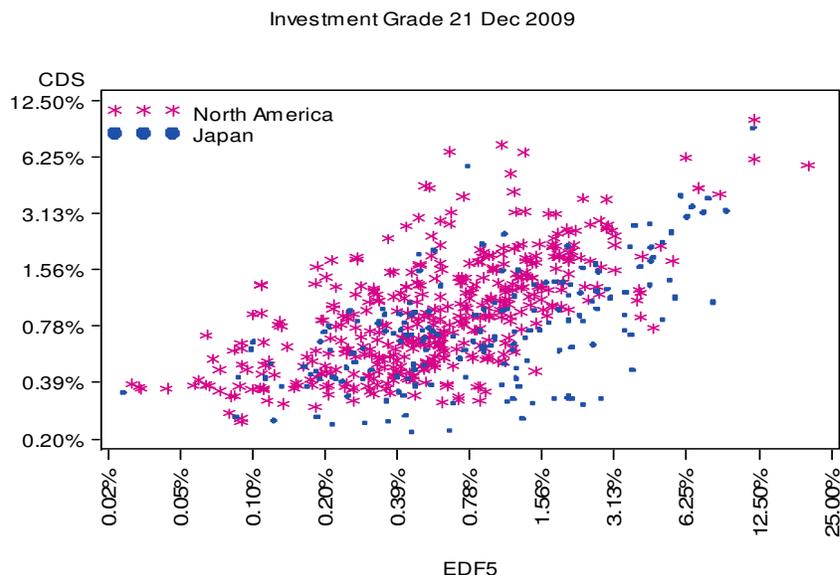


Figure 26 Japanese Spreads vs. North American Spreads

3. What market price of risk do we use for Unrated Companies?

For unrated companies, we use the high yield market price of risk. For a given EDF level, spreads for high yield (HY) companies are usually higher than spreads for investment grade (IG) companies, as shown in Figure 27.

Employing the high yield market price of risk allows the model to better match spreads for unrated companies and spreads implied by EDF levels. Institutional investors are subject to similar restrictions on unrated companies, as they are on speculative grade companies. Thus, we expect spreads to be higher for unrated companies than for investment grade companies with similar EDF credit measures. In practice, on the small sample of unrated firms with both spreads and EDF credit measures, we observe that spreads tend to be higher than those of IG firms with similar EDF credit measures.

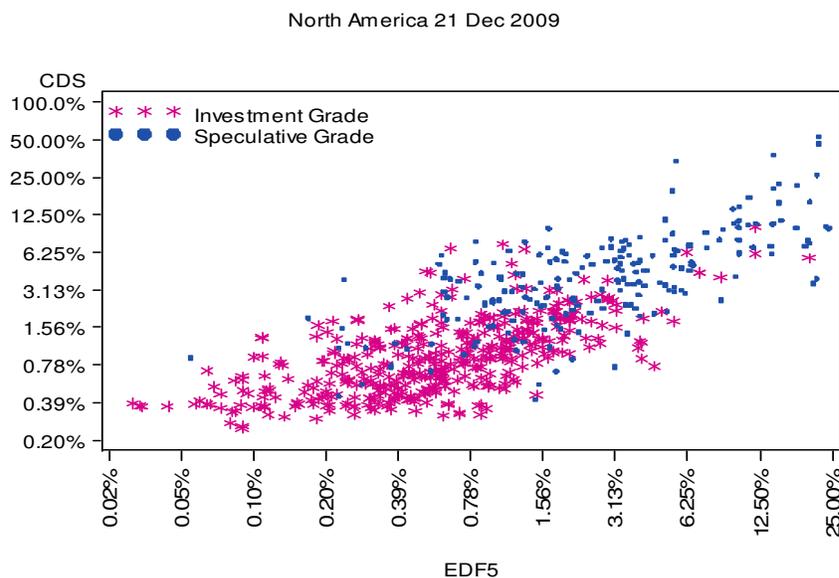


Figure 27 Spreads for HY Companies vs. Spreads for IG Companies

4. Why do we aggregate spreads across currencies and doc clauses?

CDS spreads provided by Markit are quoted separately for each entity by currency, doc clause, tier, and tenor. For each entity, we aggregate these spreads to get a single term structure of spreads for input into the Spread-implied EDF model. We see only minimal differences in spreads across different currencies. Aggregating across currencies gives us fuller coverage and more robustness.

Aggregating across doc clauses leads to fuller coverage and more stability in the model. Since spreads differ systematically by doc clause, we convert each spread into a common doc clause. The most liquid doc clause typically depends on the region and rating. However, there are some disadvantages when using the spreads of only the most liquid doc clauses. Doc clause conventions can and do change over time, and spreads of different doc clauses are on less even footing with each other. Therefore, aggregating spreads across the doc clauses enables us to limit the disadvantages and any discrepancies.

5. What are doc clauses?

The doc clause determines the types of default events that trigger payment and the admissible maturities of the deliverable bonds. In April 2009, the North American market adopted the Standard North American Contract, which calls for *no restructuring* (restructuring is not a credit event) to be the default doc clause. Nevertheless, we may continue to see a mixture of doc clauses being used in practice for different asset classes and different regions.

Broadly speaking, there are three credit events that can result in payment of protection: bankruptcy, failure to pay, and restructuring. A documentation clause, or *doc clause*, is a term clause included as part of a CDS contract. CDS contracts provide protection to debt defaults, and different doc clauses are transacted for different types of distress/default events. Currently, there are four main doc clauses:

- CR (full-restructuring)—Deliverable bonds mature within 30 years of default event.
- MM (modified-modified)—Deliverable bonds mature within five years for restructuring events and two-and-a-half years for other default events.
- MR (modified-restructuring)—Deliverable bonds mature within two-and-a-half years of default event.
- XR (no restructuring)—Restructuring is not a default event.

Due to the different levels of protection, contracts on different doc clauses are priced differently. Generally, a CR contract has the highest spread, and a XR contract has the lowest spread. Full restructuring was the original doc clause. Subsequent doc clauses were introduced after Conseco's restructuring event allowed protection buyers to profit by delivering long-dated bonds. MR became the most popular doc clause in North America and Australia. MM became the most popular doc clause in Europe. CR has remained popular in Asia as well as for sovereigns and municipals. XR tends to be used for North American high yield firms.

One of the intentions of the changes introduced in the North American CDS market in 2009 (the so-called Big Bang) was to make XR the standard doc clause (Markit Partners, 2009).²² In the case of sovereigns, economic losses to bondholders are more likely caused by restructuring than by failure to pay or bankruptcy, which likely explains why CR is a common document clause for this asset class. It also seems unlikely that the XR clause will become the dominant clause for this asset class.

6. How does the model treat different doc clauses?

When aggregating spreads across different doc clauses, we use a conversion ratio to convert clauses to a common doc clause (Full Restructuring). For the same name, different CDS spreads are quoted for different doc clauses. We convert all spreads to a common doc clause (Full Restructuring) using a multiplier. We then average over these spreads, weighting by composite depth. Thus, the output spreads are also converted to the Full Restructuring level. Converting these output spreads to the level of other doc clauses is achieved by multiplying the spreads by the doc clause ratio for the desired doc clause and region.

The multiplier is computed as follows. On each day, the ratio of spreads for CR to the other doc clause is computed for each name where both are quoted. Then, the median for this ratio is computed by region. Finally, the multiplier is the 60-day moving average of this median.

EDF implied spreads can be determined for all doc clauses using these ratios. This methodology is robust to changes in different doc clause usage over time.

7. Determining the complete term structures of Fair-value Spreads and Spread-implied EDF credit measures.

Since the relationship between EDF5 and other EDF maturities is relatively constant over time, we use static functions to derive the typical EDF term structure from EDF5.

We determine complete Fair-value Spread term structures based on the observed CDS term structures for different spread levels and regions. Recall, we can derive risk-neutral PD term structures from CDS spread term structures and vice versa using an LGD assumption. Using a 60% LGD assumption, we derive risk-neutral PD term structures

²² "The CDS Big Bang: Understanding the Changes to the Global CDS Contract and North American Conventions," March 2009, Markit Group Limited.

by region and risk level that imply spreads closest to what we observe. We assume the risk-neutral PD term structures have Weibull survival functions:

$$S(t) = \exp(-(h_0 t)^{h_1}) \quad (10)$$

With Weibull parameter h_1 varying with the cumulative 5-year risk-neutral PD Q_5 as

$$h_1(Q_5) = \exp(a + b \log(Q_5) + c \log^2(Q_5)) \quad (11)$$

The unique h_0 consistent with Q and h_1 is then

$$h_0(Q_5) = (-\log(1 - Q_5))^{(1/h_1(Q_5))} / 5. \quad (12)$$

We solve for the coefficients a , b , and c daily using a five-day window so that, given Q_5 s derived from aggregate spreads, the risk-neutral PD term structures defined by these coefficients imply spreads with the smallest relative difference to the initial aggregate spreads.

8. What are the enhancements to the CDS Fair-value Spread Calculation relative to CreditEdge Plus?

In CreditEdge Plus, Moody's Analytics provides a Fair-value Spread for both bonds and CDS. The framework in CreditEdge Plus for a FVS is conceptually comparable to the one here. However, the calculation in the new framework differs from its predecessor in several ways.

- Calibrated using CDS spreads rather than bond prices—The CDS market provides a more direct measure of credit risk than a bond price. Calibrating a market price of risk using matched CDS spreads and EDF credit measures is less involved and requires fewer assumptions than using bond prices.
- Calibrated differently for different regions—The CreditEdge Plus framework uses the same Market Price of Risk and LGD adjustment parameters for all regions. In the spread-implied EDF framework, we use different parameters for different regions (i.e., North America, Europe, Asia and South America, and Japan) and Asset Classes (i.e., investment grade versus speculative grade). We find that we can obtain a more accurate calibration by using different parameters for different regions.
- Term Structure of FVS based on typical term structure of CDS spreads—The framework in CreditEdge Plus derives a term structure of FVS by using the EDF term structure and the relationship between a risk-neutral term structure and a physical term structure implied by option pricing theory. For spread-implied EDF, we use the typical term structure of CDS spreads found in the market at that time to derive the term structure of FVS. We find this approach yields term structures of FVS that are more consistent with the term structure of actual spreads at different stages in the credit cycle.
- Fair-value Spread is based on Full Analytics—The conversion of an EDF into a FVS in CreditEdge Plus was based on a zero coupon bond approximation to a bond or CDS contract that pays a coupon for a specified time period.²³ The zero coupon approximation is a good approximation when credit risk is low, but it begins to break down when credit risk is elevated. The spread-implied EDF model replaces this approximation by accounting for all cash flows as described in Section 3.3. As a result, the FVS is accurate even when credit risk is very high.

²³ See Arora, Agrawal, and Bohn 2004.

Acknowledgements

We thank David Munves, Vlada Stein, and Zhao Sun for their helpful comments and suggestions. We appreciate Christopher Crossen for his editorial assistance.

Copyright © 2010 Moody's Analytics, Inc. and/or its licensors and affiliates. All rights reserved.

References

- Agrawal, D., N. Arora and J. Bohn. 2004. "Parsimony in Practice: An EDF-Based Model of Credit Spreads," Moody's KMV White Paper.
- Arora, Navneet, Jeffery Bohn, and Irina Korablev. 2005a. "Power and Level Validation of the EDF™ Credit Measure in the U.S. Market," Moody's KMV.
- Basel Committee on Banking Supervision, "Strengthening the Resilience of the Banking Sector," issued for comment December 2009, Bank of International Settlements.
- Black, Fischer, and J. Cox. 1976. "Valuing Corporate Securities and Liabilities: Some Effects of Bond Indenture Provisions," *Journal of Finance* 31: pp 351–367.
- Black, Fischer and Myron Scholes. 1973. "The Pricing of Options and Corporate Liabilities." *Journal of Political Economy*, Vol. 81, No. 3, pp 637-654.
- Cantor, Richard and Chris Mann. September 2006 "Analyzing the Tradeoff Between Ratings Accuracy and Stability," Moody's Investors Service.
- Crosbie, Peter and Jeffrey Bohn. 2003. "Modeling Default Risk," Moody's KMV White Paper.
- Dwyer, Douglas and Irina Korablev. 2007. "Power and Level Validation of Moody's KMV EDF™ Credit Measures in North America, Europe, and Asia," Moody's KMV White Paper.
- Dwyer, Douglas and Shisheng Qu. 2007. "EDF™ 8.0 Model Enhancements," Moody's KMV White Paper.
- Gisiger, Nicolas. 2009. "Risk-Neutral Probabilities Explained".
- Gray, Dale and Samuel Malone. 2008. *Macrofinancial Risk Analysis*, John Wiley and Sons.
- Kealhofer, Stephen. 2003a. "Quantifying Credit Risk I: Default Prediction." *Financial Analysts Journal*, January/February 2003, 30–44.
- Kealhofer, S., "Quantifying Credit Risk II: Debt Valuation." *Financial Analysts Journal*, May/June 2003, 78-93.
- Korablev, Irina and Shisheng Qu. 2009 "Validating the Public EDF™ Model Performance During the Credit Crisis," Moody's KMV White Paper.
- Merton, Robert C. 1974. "On the Pricing of Corporate Debt: The Risk Structure of Interest Rates." *Journal of Finance*, Vol 29, May, 449–470.
- Merton, Robert. 1992. *Continuous Time Finance*, Blackwell Publishing.
- Munves, David and David Hamilton. December 2007 "Moody's Market Implied Ratings: Description, Methodology, and Analytical Applications," *Moody's ViewPoints*.
- Schaefer, Stephen, M. and Ilya A. Strebulaev. October 2008. "Structural Models of Credit Risk Are Useful: Evidence from Hedge Ratios on Corporate Bonds," *Journal of Financial Economics*, Volume 90, Issue 1, pp 1-19.
- Zhang, Jing and Zan Li. 2009. "Investing in Corporate Credit Using Quantitative Tools," Moody's KMV.