MOODY'S ANALYTICS

MODELING METHODOLOGY

Credit Risk Modeling of Public Firms: EDF9

Authors
Pooya Nazeran
Douglas Dwyer

Acknowledgements
We are grateful for the advice and insights provided by Vera Arsova, Clara Bernard, Louis Chapman, Ramkumar Chandramohan, Nan Chen, Christopher Crossen, Houman Dehghan, Min Ding, Jian Du, James Edwards, Danielle Ferry, David Hamilton, Fangze Huang, Glenn Levine, Zaid Marji, Jacob Matsibekker, Mathew McDonald, David Munves, Paolo Persurich, Cristina Pieretti, Tanya Roosta, Amrit Sandhir, Sukie Sun, Zhao Sun, Sunny Wong, Pierre Xu, Sherry Yang, Jing Zhang, Sue Zhang, and Yaran Zhang.

Contact Us

Americas
+1.212.553.1653
clientservices@moodys.com

Europe
+44.20.7772.5454
clientservices.emea@moodys.com

Asia-Pacific (Excluding Japan)
+85 2 3551 3077
clientservices.asia@moodys.com

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

Abstract

EDF9 — the 9th generation of the Moody's Analytics Public Firm EDF™ (Expected Default Frequency) model — expands the frontiers of structural credit risk modeling. EDF metrics are forward-looking probabilities of default, available on a daily basis for 35,000-plus corporate and financial firms. The updated EDF9 model incorporates insights attained by evaluating the behavior of the prior version, EDF8, over the course of the recent financial and sovereign debt crises. EDF9 also utilizes a larger dataset, given the global expansion of the equity markets, enhanced data quality, and improvements in computational performance. This paper explains the Public Firm EDF model's basic methodology, as implemented in the new EDF9 version. It also explains the structure and conceptual grounding of the core model and highlights new model features. Additionally, we provide empirical details of the model's calibration.

Table of Contents

1. Introduction

Measuring a firm's probability of default (PD) is one of the central problems of credit risk analysis. Moody's Analytics' Public Firm EDF (Expected Default Frequency) model has been the industry-leading PD model since its introduction in the early 1990s. Since that time, the model has undergone considerable development, and it continues to evolve, while providing unequalled global coverage of public firms on a daily basis.

Since the release of the previous version, EDF8, the global economy experienced a financial crisis that expanded into what has been called the Great Recession.[1] Europe dealt with sovereign credit crises. Monetary policy took on an unprecedented role in the broader economy. Considerations of quantitative easing, as well as bailouts, became important pricing factors for fixed-income securities. The term SIFI[2] was introduced, and new regulations were introduced around the world, affecting capital market structures and the roles both private and public institutions play. At the same time, the BRIC[3] countries and other emerging markets continued to expand, resulting in a large increase in the number of publicly traded firms. Subsequently, data feeds expanded coverage outside of North America.

In light of these events, we revisited our process of measuring PD's. The updated model and data used in EDF9 reflect new lessons learned from the most recent credit cycle. An important revision expands the EDF9 dataset for financial firm failures, which improves our ability to predict defaults for these firms. Financial firms typically present a different default dynamic when compared to non-financial corporations. The increase in the number of financial firm failures in our default database enabled us to refine the Public Firm EDF model's treatment of financial firms in multiple ways. EDF9 also benefits from the global growth of exchanges. This expansion enables more granular modeling of certain risk-factors, which tailors the risk metric to the specific country where a firm resides.

What's New in EDF9

While EDF9 uses the same conceptual framework as EDF8, we make a number of important enhancements.

We have updated and improved DD-to-EDF measure mapping. The mapping for non-financial firms now incorporates international defaults from 1987 – 2014.

For the first time with the Public Firm model, we introduce a financial firm-specific DD-to-EDF mapping that uses international financial firm failures spanning 1987 – 2014. EDF values for financial firms represent the risk of either a default or a government bailout. The numerous financial firm failures observed during the last seven years makes such a mapping feasible. Financial firm-specific mapping generates more stable EDF values over time for this sector. Consequently, EDF9 values are typically lower than EDF8 for the riskiest financial firms during a bad stage of the cycle, and, conversely, EDF9 values are typically higher than EDF8 values for the safest financial firms during a good stage of the cycle. We also revisited the calculation of the financial firm default point. Relative to EDF8, we employ a relatively simple and parsimonious calculation. This change makes the model more transparent for financial firms.

We also made refinements to the volatility calculation. These refinements fine-tune a volatility measure that is now implemented consistently across the globe, and granularity is now country-level. Additionally, the volatility measure better reflects the volatility trends of the specific country and industry that a firm operates within. This volatility measure is somewhat higher when volatility is low on a historical basis.

Finally, we produce EDF values back to 1963 using one uniform algorithm.

EDF9 retains the principal structure and conceptual grounding of the core model. The model remains an option-pricing based, structural model. We continue to utilize Vasicek's formulation of option contracts, as implemented in the original KMV model. EDF9 also continues to summarize all credit risk-related information into one metric, called Distance-to-Default (DD). We then map the DD to the probability space to obtain the EDF credit metric.

This paper explains the recent updates and changes made to the model, the rationale for the changes, and their implications for the one-year EDF measure.[4]

[1] We use the term "Great Recession" to refer to the period that began with a financial crisis in late 2007 and then spread, impacting non-financial firms, the unemployment rate, other macroeconomic variables, and sovereign credit risk. This recessionary period ended in mid-2009.

[2] Systematically Important Financial Institutions.

[3] Brazil, Russia, India, and China.

[4] Chen, et. al (2015) provides an overview of many aspects of the EDF9 framework. By focusing on the one-year EDF credit measure, this paper provides an in-depth discussion of Section 4 in that overview paper.

We organize the paper as follows:

Section 2 presents the theoretical background of the Public Firm EDF model and discusses its basic structure and components. For readers new to structural modeling, Appendix B explores the model from a perspective more familiar to those with fundamental analysis background.

Section 3 details specific EDF9 model updates and changes, as well as practical implementations of the model.

Section 4 discusses model changes pertaining to the calibration of the model's DD-to-EDF transformation.

Section 5 provides general measures of model performance.

Section 6 summarizes and concludes.

Appendix A describes the products, services, and models related to the Moody's Analytics Public EDF model suite. These products include: CreditEdge®, RiskCalc™, GCorr™ and RiskFrontier™.

2. Theoretical Background of the Public Firm Model

This section presents the Public Firm EDF model's background and discusses its basic structure and components.

While many readers are comfortable with the underlying theoretical framework of the KMV structural model, others may be new to it. Many authors provide useful descriptions of how the approach works using option-pricing theory. For in-depth overviews, please refer to Crosbie and Bohn (2003), Ranson (2005), Caouette, et.al (2008), Bohn and Stein (2009), Duffie and Singleton (2012). For readers new to the model, Appendix B provides a more conceptual walk-through that describes the motivation behind the framework by transforming an accounting-based measure of leverage into a heuristic approximation to the Distance-to-Default concept.

The Public Firm EDF model belongs to the class of structural credit risk models, pioneered by Fischer Black, Robert Merton, and Myron Scholes. This model, originated by KMV,[5] takes an option-pricing based approach to credit risk. Only when an option contract expires in-the-money, does it receive a payoff. Naturally, the valuation reflects the probability of the contract expiring out-of-the-money. Due to limited liability laws, the market value of a publicly traded company's equity is lower bounded at zero, giving it call option-like characteristics. Hence, the default probability is embedded in the stock price. The stock price, however, contains other information as well. This complexity makes extracting the PD an interesting problem. Option pricing research provides a mathematical structure, within which, we can study default probabilities.

Stochastic calculus, in the context of Merton's model, can be used to understand the link between the market value of assets and the market value of equity. Let the market value of a firm's underlying assets follow a stochastic process:

$$\boldsymbol{dA = \mu_A A\, dt + \sigma_A A\, dw} \tag{1}$$

where

$\boldsymbol{A, dA}$ are the value of firm's assets, and their change.
μ_A, σ_A are two parameters governing, respectively, the drift and risk of assets.
$\boldsymbol{dw}$ is the Wiener process

Theoretically, default occurs when the firm's asset value falls to the default point level. The probability that asset values drop to such levels depends on the distribution of asset returns, which is, in part, characterized by asset volatility. We can formulate this by positing the following as the question of interest:

$$\boldsymbol{PD_t = Pr[A_t \leq D] = Pr[ln(A_t) \leq ln(D)]} \tag{2}$$

where $\boldsymbol{D}$ is the default point, and the subscript $\boldsymbol{t}$ indicates time in the future. By Ito's lemma, the equation describing the asset value process implies:

$$\boldsymbol{ln(A_t) = ln(A_0) + \left(\mu_A - \frac{1}{2}\sigma_A^2\right) \cdot t + \sigma_A \sqrt{t} \cdot w} \tag{3}$$

[5] KMV was acquired by Moody's Corporation in 2002. KMV is currently part of Moody's Analytics, a subsidiary of Moody's Corporation.

Where $\boldsymbol{w}$ is the source of stochasticity, with a known distribution. In that case, the probability equation can be written as:

$$\boldsymbol{PD_t = Pr\left[ln(A_0) + \left(\mu_A - \frac{1}{2}\sigma_A^2\right)\cdot t + \sigma_A\sqrt{t}\cdot w \leq ln(D)\right]} \quad (4)$$

which can then be simplified to:

$$\boldsymbol{PD_t = Pr\left[\frac{ln\left(\frac{A_0}{D}\right)+\left(\mu_A-\frac{1}{2}\sigma_A^2\right)\cdot t}{\sigma_A\sqrt{t}} \leq -w\right]} \quad (5)$$

The factor inside this formula is Distance-to-Default. In the finance literature, we find derivations for a heuristic version, presented on the left-hand side of equation 6.[6] Appendix B shows an alternative derivation of equation 6, from a balance sheet perspective.

$$\boldsymbol{\frac{A-D}{A\cdot\sigma_A} \approx \frac{ln\left(\frac{A}{D}\right)}{\sigma_A}} \quad (6)$$

In order to compute DD, we must estimate A_0 and σ_A. Before discussing our implementation of the estimation process, we discuss how to estimate A_0 and σ_A using the Black-Scholes formula.

As a special case of the Merton model, it is possible to use the Black-Scholes formula to illustrate the process behind estimating asset value. In practice, we use the Vasicek-Kealhofer (VK) model to calculate EDF measures, and we discuss the particular implications of that to asset value calculation later on in this paper.

Equity is the present value of earnings, the residuals of cash flows, after all others obligations have been met. But equity is subject to limited liability. This relationship is analogous to a call option, which has a claim to the assets in excess of the strike price. In our case, the strike price is the book value of liabilities.

We can exploit this characteristic of equity to asset value in a Black-Scholes context:

$$\boldsymbol{E = A\,\Phi(d_1) - e^{-rT}\,X\,\Phi(d_2)} \quad (7)$$

where

E is the equity value.

$d_1 = \frac{ln\left(\frac{A_0}{D}\right)+\left(r+\frac{1}{2}\sigma_A^2\right)\cdot T}{\sigma_A\sqrt{T}}$

$d_2 = \frac{ln\left(\frac{A_0}{D}\right)+\left(r-\frac{1}{2}\sigma_A^2\right)\cdot T}{\sigma_A\sqrt{T}}$

r is the risk-free interest rate

Φ is the cumulative standard normal distribution function.

While in an option-pricing context A is taken as given and E is calculated, for our purposes, we can use the observed E to estimate the market value of assets.

$$Asset = Call\ Option + Strike\ Price - Put\ Option$$
$$Asset = Equity + Liabilities - Default\ Risk\ Premium$$

While this estimation of asset value reflects the default risk premium,[7] equation 7 is not practical as, for instance, one cannot reflect the value of preferred stock in said premium, different maturities of liabilities cannot be specified, and equity cannot be treated as a perpetual option without an expiration date. We discuss these restrictions and our solution for them in Section 3.

To estimate σ_A, we need to understand how changes in equity value are related to changes in asset value. As the residual of asset value after obligations are met, equity value's dependence on changes in asset value can be described with partial derivatives.

$$\boldsymbol{dE = \frac{\partial E}{\partial A}\cdot dA} \quad (8)$$

[6] For instance, refer to Crosbie and Bohn (2003).

[7] For the purpose of this discussion, the default risk premium is not in yield space, but rather in price space. In other words, it can be thought of as the price differential between the market value of a default-risky bond and the market value of a bond with identical cash flows, but without default-risk.

$$\frac{dE}{E} = \frac{\partial E}{\partial A} \cdot \frac{A}{E} \cdot \frac{dA}{A} \tag{9}$$

$$r_E = \frac{\partial E}{\partial A} \cdot \frac{A}{E} \cdot r_A$$

There are two key terms governing the relation between asset returns and equity returns, notably $\frac{A}{E}$ and $\frac{\partial E}{\partial A}$. To illustrate the role of the first term, consider an institution with $1B in assets and zero liabilities. Each change in asset value propagates to equity valuation dollar-for-dollar. Hence, the company's assets and equity will have the same volatility. In the case of a firm with $1B in assets but $500MM in liabilities, a $10MM change in assets changes asset values by 1%, while the effect of a $10MM change in equity is 2%.

However, that change is not exactly 2% because of $\frac{\partial E}{\partial A}$. Take a bank for example, with $1B in assets and $900MM in liabilities. Without limited liability, $\frac{\partial E}{\partial A}$ would be one, and a 1% change in assets results in an 11% change in equity. But, given limited liability, as the value of assets becomes closer to the value of total liabilities (strike price), the derivative becomes smaller, reflecting increased default risk. Consequently, the $10MM is distributed between both changes in the market value of liabilities and changes in equity valuation, with a change in equity of less than 11%.

The $\frac{\partial E}{\partial A}$ reflects the effect of the optionality in equity and is what option investors call delta. The delta parameter appears in the application of Ito's lemma, which provides a more formal derivation of the relation between asset volatility and equity volatility. We end this section with the formal derivation.

Since the value of equity is determined, in part, based on the value of assets, we can describe it in a simple notation as

$$E = f(A)$$

By Ito's lemma, we arrive at:

$$\boldsymbol{dE = \left[\frac{\partial E}{\partial t} + \frac{\partial E}{\partial A}\mu_A A + \frac{1}{2}\frac{\partial^2 E}{\partial A^2}\sigma_A{}^2 A^2\right] dt + \frac{\partial E}{\partial A}\sigma_A A\, dw} \tag{10}$$

On the other hand, dE itself is a stochastic process and can be described as:

$$\boldsymbol{dE = \mu_{E,t} E\, dt + \sigma_{E,t} E\, dw} \tag{11}$$

Matching the second term of each equation, we arrive at:

$$\boldsymbol{\sigma_{E,t} = \frac{\partial E}{\partial A} \cdot \frac{A}{E} \cdot \sigma_A} \tag{12}$$

This formula implies that equity volatility changes as the leverage of a firm changes. Therefore, one cannot use it to determine asset volatility from equity volatility if equity volatility is measured as the standard deviation of equity returns over a window where leverage changes.

3. Practical Implementation of the VK / EDF Model

This section covers implementation considerations of the three primary drivers of the model: asset value, default point, and asset volatility. Figure 1 provides a visual representation of the relationship between these components of the model and their dependencies. We also highlight the significant improvements found in EDF9.

The VK model, the foundation for the Public Firm EDF model and EDF9, extends on the Merton model in a number of ways. The extensions are intended to make the Merton model more useful in practice. The extensions allow the Merton model to be applied to firms with different types of capital structures. The VK model includes both short-term and long-term debt. Further, it allows for convertible securities and cash leakages. Finally, it incorporates a concept of preferred stock.

In EDF9, we make the following changes: For all firms, an adjustment is made to the default point to reflect the cost of capital. For financial firms, we simplify the specification of the default point. Asset volatility is the combination of empirical volatility and modeled volatility. For empirical volatility, we now measure it consistently throughout the world using a three-year window of

weekly returns. For modeled volatility, we incorporate higher granularity with respect to geography and industry. Last, a new component has been introduced into the asset volatility calculation, which makes it more forward-looking.

Figure 1 Public EDF Model and Drivers

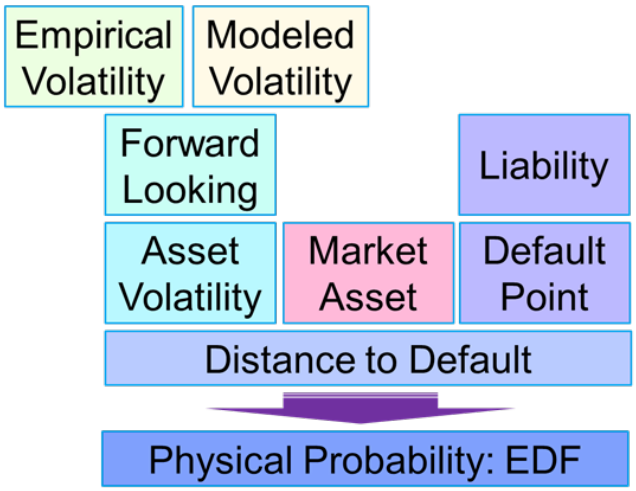

3.1 Market Value of Assets

In the VK model, the market value of a firm's assets equals the sum of the market values of common stock, convertible securities, short-term debt, and long-term debt. We begin by describing the concept of asset value and then discuss each component in turn.

The concept that balance sheets' balance states is that the value of assets equals the value of liabilities and equity combined. While the balance sheet reflects their book valuations, the equality continues to hold for market valuations. With the market value of publicly traded equity, this implies we can calculate the market value of assets if we know the market value of liabilities.

The book value of debt (money a firm has borrowed from a lender) is typically the discounted value of the remaining future payments, discounted at whatever interest rate the firm and the lender agreed to when the debt was originated. As both the risk-free rate and the firm's credit risk change over time, the book value of a debt typically differs from its market value — how much a third party would pay the firm's creditor for the debt. Had the firm's liabilities been risk-free, the market value of liabilities could have been calculated by discounting the future payments with a risk-free interest rate. In practice, because of default risk, the market value is less.[8]

$$\boldsymbol{Market\ Value\ of\ Liabilities = Riskless\ Value\ of\ Liabilities - Default\ Risk\ Premium} \quad (13)$$

As a firm's liabilities are complex, it is typically not possible to directly observe the market value of a firm's liabilities. Oldrich Vasicek developed a model in the 1980s that estimates both the market value of assets and the market value of liabilities from the market value of equity. The model takes asset value, asset volatility, the risk-free rate, as well as a few balance sheet and income statement items and returns the difference between the risk-free and risky value of the debt, namely the default risk premium.

The standard Merton model assumes two types of claims to the cash flow generated by a firm — debt with no coupons and equity with no dividends. This model posits that the underlying assets can be represented by a geometric Brownian motion process that is parameterized by volatility and a drift term. The debt is a one-time payment at a specified point in time. In this context, the actual solution is the Black-Scholes option-pricing formula.

The structural model that forms the foundation for the Public Firm EDF model is the Vasicek-Kealhofer (VK) model. The VK framework assumes five types of claims on the firm's cash flows: short-term liabilities, long-term liabilities, convertible securities, preferred stock, and common stock. Incorporating these different types of contingent claims into the model changes both the asset value formula and the boundary conditions.

[8] One can think of the default risk premium as the present value of a CDS contract. Indeed, it is for this risk that credit default swap (CDS) contracts exist and trade. If we had the present value of a CDS contract, in other words, the upfront fee demanded on a hypothetical zero-coupon CDS contract, we could calculate market value of liabilities. For most publicly traded firms, however, there is no CDS contract, and, where there is one, it may not be liquid, and, where it is liquid, it may only reflect the risk premium on a subclass of the firm's liabilities.

3.1.1 Dividends Coupons and Interest Expense

Cash leakages in the forms of dividends on stock, coupons on bonds, and interest expense on loans impact both default probabilities and debt value. For example, consider two firms with identical assets and debt, but one pays a larger dividend. The firm paying the larger dividend has a higher default probability, even though the dividend may be cut in the event of distress; any higher dividend payments made until distress became apparent would reduce the cash flow available to service debt. The VK model incorporates cash outflows directly into the different types of claims on a firm's assets.

3.1.2 Convertible Securities

Companies may issue securities that can be converted into equity at a specified conversion rate. Such securities are often preferred stock, but bonds can be convertible as well. By issuing such securities, the firm is effectively selling a portion of the upside return that otherwise belongs to the common stock holders. Consider two firms: A and B. Both have the same assets and debt. Company B has a convertible security outstanding, so the fully-diluted shares outstanding exceeds the common shares outstanding. Under this scenario, Company B has a lower market value of equity than Company A, even though the default probability is the same. Holders of common stock in company B sell a portion of the upside return to the holders of the convertible security. This difference becomes reflected in one of the boundary conditions for the VK model — that, as the asset value of the firm becomes arbitrarily large, the derivative of equity value, with respect to asset value, becomes equal to the ratio of the shares outstanding, divided by the number of fully diluted shares outstanding. The dilution effect of convertible securities reduces the sensitivity of the value of equity to the value of assets. If this dilution effect is ignored, when observing the equity value of a company that has a large number of convertible securities, one underestimates the market value of assets and overstates the probability of default.

3.1.3 Current and Long-term Liabilities

In structural models that use an absorbing default barrier, two approaches have been taken to define the barrier. The default event can be driven by creditors forcing the company into default when the asset value falls below the barrier. Alternatively, if there are no protective covenants, the company can choose to default when the value of equity falls to zero — if the value of assets falls below a certain threshold, then the equity holders choose to stop making payments on the debt and, thereby, turn over the firm's assets to debt holders.[9] In the VK model, the barrier is exogenous, in the sense, that creditors put the debt back to the firm as soon as the value of assets hits the value of the default point. This barrier forms the second boundary condition of the model — the value of equity equals zero when the value of the firm's assets equals the default point.

3.1.4 Preferred Stock

Preferred stock has both equity and debt characteristics, and the model takes both these aspects into consideration. The VK model can incorporate various types of preferred stock, including tradable preferred stock and convertible preferred stock.

3.1.5 Implementation

Vasicek's formula poses two implementation challenges: its dependence on asset value and its dependence on asset volatility.

We require the default risk premium to estimate the market value of liabilities in order to subsequently calculate the market value of assets. On the other hand, we need the market value of assets for the default risk premium calculation. We resolve this issue with the put-call parity equation, directly deduced from the definition of option contracts. In our context, it implies

$$\boldsymbol{Asset\ Value = Equity\ Value + Riskless\ Value\ of\ Liabilities - Default\ Risk\ Premium} \quad (14)$$

Exposure to default risk is similar to shorting a put option, as it makes the creditor subject to the risk of assets dropping below the default point. By solving this equation, we obtain estimations for asset value, put option, and hedge ratio.

Asset volatility is also a necessary input, which we address in later sections. Before proceeding, it is important to clarify that we do not change the put option equation in EDF9. Nonetheless, as inputs such as asset volatility have typically changed, the resulting asset value for a specific name at a specific point in time is typically different when compared to EDF8.

[9] For early implementations of so-called exogenous and endogenous default barriers, see Black, Cox (1975) and Leland (1994), respectively.

3.2 Default Point

What's New in EDF9: *For all firms, we introduce a cost-of-capital adjustment to the default point. Furthermore, we reformulate the default point for financial firms.*

The long-standing practice of defining the default point has been "current liability + ½ long term liability." This formula originates from KMV research and has been employed in the academic literature. As part of the EDF9 research, we reevaluated the formula against many alternative possibilities. We found that the other approaches were not consistently better than the long-standing practice for non-financial firms. For financial-firms, however, we make a change to the default point calculation, which we discuss in the next section.

Looking at model behavior in different interest rate environments, we also add a cost-of-capital adjustment to the default point, for all business types. In a high interest rate environment, a firm's business partners are less willing to rollover credit and allow accumulation of unpaid bills, when compared to a low interest rate environment. Since, in a high interest rate economy, working capital depletes faster and access to credit becomes expensive, all else being equal, the default rate is expected to be higher. Controlling for all other aspects of credit risk, we indeed find that an unadjusted default point underestimates credit risk during the 1980s and 1990s and overestimates during the more recent low-rate periods. Adjusting the default point by the cost-of-capital, using 1+r, we find improvements to the EDF measure across the different interest rate environments. As rates have been fairly low during 2014 and 2015, to the extent that they increase in the future, all else being equal, the model should estimate higher EDF measures.

3.2.1 Financial Firms

Estimating the default point as current liabilities plus one-half long term liabilities for financials is problematic, as financial firms typically do not report current liabilities as a separate line item on their balance sheets.

While a financial firm's balance sheet lists long-term debt, identification of the term structure of the remaining balance sheet is not so straightforward. For the Public EDF model, we had developed a method for estimating the long-term component of the remaining liabilities, using a firm's size, as well as sector-level information. However, in EDF8 and earlier versions of the model, we also faced another challenge — scarcity of default events among publicly traded financial firms at the time. This issue further complicated the modeling of a default point for such firms.

Following the recent financial crisis, we possessed sizeable financial firm default data, which helps us understand their credit risk and default dynamics much better.

Having confirmed that the long-term liabilities of a non-financial corporate imposes half the burden of current liabilities, we hypothesized that long-term debt also plays a different role in a financial firm's default point. We further separated deposits from the rest of liabilities. In other words, we separated total liabilities of a financial firm into three components: long-term debt, deposits, and everything else. In searching for differential impacts of different liability types, we empirically evaluated the data. We assessed the performance of the model for a variety of different default point definitions. We were unable to find a consistent and statistically meaningful difference between these three components of total liabilities.

Our findings can be attributed to the sophistication of the instruments on a financial firm's balance sheet. The same sophistication that makes it difficult for most accountants to identify and separate current liabilities from long-term liabilities could affect the informativeness of the balance sheet. Furthermore, relative to non-financial corporates, financial firms can restructure their liabilities and change their tenor more flexibly. Naturally, all these possibilities comprise to make a financial balance sheet more opaque and its credit risk harder to measure.

Based on our empirical findings, we set the default point as a percentage of total adjusted liabilities. We calibrate that percentage as 75%. The mentioned adjustment to total liabilities excludes minority interest and deferred tax, as they are not sources of credit risk.

A side benefit of this default point definition is more stable EDF measures for financials, relative to EDF8. In EDF8, the financial firm default point depended on long-term debt, but on occasions, this value tended to vary dramatically between quarterly and annual statements.[10] Subsequently, this variation introduced spurious changes in EDF8's default point for financial firms, but does not have any effect on EDF9's default point.

[10] For an explanation on the role of variation of debt value on default point, refer to Knowledge Base FAQ 388 (2006).

3.2.2 Off-Balance Sheet Items

The default point reflects the current and total liabilities of the balance sheet. A common concern is the impact of off-balance sheet items on the EDF measure. However, EDF measures' sensitivity to off-balance sheet items is low, as the model compensates in its estimation of volatility.

For a firm with off-balance sheet items that are a source of credit risk, total liabilities is underreported. Subsequently, the asset value and the default point are understated. Take, for example, Bank B with equity valued at $65B,[11] and $1T[12] in total liabilities. We estimate asset volatility at approximately 3% and asset value of $1005B. The $5B difference between the asset value and the total liabilities value tells us the estimated default risk premium equals $60B.

Let us assume 10% of the liabilities are off-balance sheet and are not reflected in the $1T value. This means we underestimate the default point by about 10%, and, instead of $750B,[13] it should be $825B. On the one hand, by the balance sheet principle, the asset value should be $75B more. On the other hand, this leverage increase raises the default risk premium from $60B to about $67B, as the higher-leverage means higher default risk. All in all, asset value increases by approximately 9%, while the default point increases by about 10%. Our measure of leverage, asset value divided by default point, therefore, decreases by 1%, reducing DD by about 1%.

But this result affects our estimate of asset risk as well. If true total liabilities is $1.1T and true leverage 1% different, that means we de-levered equity returns with a ratio off by1%. Now, if we recalculate asset volatility with these corrected leverage ratios (further discussed in Section 7), instead of 3%, we find asset volatility is about 2.97%.

Since asset value / default point is the primary driver of the DD's numerator, and in the denominator it is asset volatility, the 1% reduction in the numerator is offset by the 1% reduction in the denominator, keeping the ratio by-and-large the same. We consistently observe that these forces offset one another in a way that the model produces default probabilities that reflect the credit risk of the obligor, as if all sources of risk had been reported on the balance sheet, and so long as they do not constitute a substantial portion of the liabilities.

All-in-all, as long as the market appreciates the existence of off-balance sheet items and reflects them in levering-up asset returns to equity returns, the model reflects their risk in its calculated EDF measure via the interaction of asset volatility and leverage. This offsetting effect demonstrates the important role played by the asset volatility estimation process.

3.3 Asset Volatility

Asset volatility serves two purposes in the model: one is in constructing the DD, and the other is in calculating the put option. Measuring asset volatility, however, introduces multiple challenges. Unlike market capitalization, neither equity volatility nor asset volatility is observable. These statistics can only be calculated over a window of returns, which subsequently introduces two challenges: one emerges for new firms, as sufficient history is not available, and the other in the backward-looking characteristic of such a measure.

We measure the volatility of assets over a three-year window, using data at weekly frequency from 156 observations. In addition to empirically measuring asset volatility, we model it as well, to account for new firms where sufficient history is not available. For each firm, modeled volatility is the asset volatility of the typical firm with similar characteristics. We find three characteristics informative of asset risk: size, location, and business type. By taking a weighted-average of empirical volatility and modeled volatility, we compensate for sampling variability of the limited history of asset returns for any one firm.[14]

The Public EDF model forecasts the default probability over the next year, and hence, requires the asset volatility over the next year, which is not available. Even with a full window of observations, the current asset volatility is not observable either, but we have asset volatility during the past three years. The modeled volatility is calibrated to empirical volatility and, hence, does not change its temporal features. Having observed the cyclical behavior of aggregate asset risk, at both the country-level and industry-level, we adjust the estimated asset volatility to become more prudent for the future. In the next few sections, we provide more details on empirical and modeled volatility, as well as the forward-looking adjustment.

[11] Sixty-five billion dollars

[12] One trillion dollars

[13] As mentioned in the previous section, the default point is 75% of adjusted total liabilities. For the purpose of this example, we assume deferred tax and minority interest are zero; so total adjusted liabilities is the same as total liabilities.

[14] For the intellectual origins of combining empirical and modeled asset volatility, see Vasicek (1973).

3.3.1 Empirical Volatility

What's New in EDF9: *We compute empirical volatility using a three-year window of weekly returns globally. EDF8, in contrast, used a five-year window of monthly returns for some regions.*

The process of empirically calculating asset volatility begins with estimating asset returns as de-levered equity returns. We then use asset returns during the three most recent years for the estimation.[15] In this section, we explain how we estimate asset returns, how we deal with illiquidity, and how corporate events affect estimation.

The market value of assets equals the sum of the market value of equities plus the market value of liabilities. In order to compute asset returns, we must take into account how changes in the price of equity and the price of liabilities impact the market value of assets. After controlling for shares outstanding, as the stock price increases, the market value of assets increases. As interest rates increase, the market value of liabilities declines, with the impact increasing in the tenor of the debt. Returns on long-term debt are estimated to a first-order approximation using the concept of Macaulay duration. We then compute asset returns as the sum of de-levered equity returns and the returns on long-term debt multiplied by the ratio of the long-term debt to asset value.

Consistent with earlier versions of the Public EDF model, we apply an iterative process. The process begins with a guess for asset volatility and subsequent calculation of asset value and hedge ratio. Hedge ratio is a byproduct of the asset value calculation process. Having calculated a time series of asset returns, we use it to update our asset volatility estimation. If the resulting volatility has converged to the starting volatility, we consider the asset return final, otherwise, we iterate, taking this new volatility as the starting one.

The iterative approach is superior to "solving for two equations and two unknowns" found in some academic and commercial implementations of structural models. Since equity volatility cannot be measured instantaneously and depends on a window, implicit in de-levering equity volatility is the assumption that the leverage ratio, as well as hedge ratio, have been constant throughout the window. As equity volatility changes with the market capitalization of the firm, this assumption is not valid in practice

If a firm's stock is not liquid enough to have weekly returns, during the period of its illiquidity, the model resorts to using data at the monthly frequency. To compensate for sampling variability resulting from fewer observations, we expand the three-year window. For the extreme case where a firm only has monthly returns available, we use a five year window.

Major corporate events such as mergers and acquisitions introduce another challenge. If such activity is about to occur or has even happened but is not yet reflected in the new financial statements, the model is not aware of them, whereas, the market participants take the new capital structure into consideration when deciding to buy or sell at the prevailing price. Consequently, market prices and equity returns reflect such information. When the market knows the current financial statements are outdated, our de-levering process results in asset returns too volatile when total liabilities is too small (e.g., when there is M&A activity) and asset returns that understate volatility when total liabilities is too large (e.g., in the case of a spin-off). To account for this issue, the model has an event-adjustment mechanism that retroactively adjusts the asset returns of the event period in order to correctly estimate asset volatility. While still somewhat relevant, pre-event asset returns are not as relevant to post-event asset risk. So we lower the weight on pre-event asset returns. When empirically calculating volatility, we use the magnitude of the event to determine the weight on pre-event observations.

To further improve our asset volatility estimation, the model has a mechanism for identifying and winsorising outliers. Outlying asset returns can exist for a variety of reasons, but are not helpful in the estimation of asset risk, particularly given the model's need for next year's risk. Such winsorising of outliers reduces spurious changes in volatility and, hence, EDF measures.

3.3.2 Modeled Volatility

What's New in EDF9: *We increase the modeling granularity with respect to business location and business type.*

In general, modeled volatility is used to compensate for sampling error. A new firm does not have any equity returns for empirical measurement. Hence, the model relies 100% on modeled volatility. Over time, as more and more observations become available, the model continuously shifts to place more weight on empirical volatility. At the extreme, where 156 observations are available, the role of modeled volatility reduces to about 20%, but never 0%. Just as a Bayesian-prior is used to improve the estimation of a statistic, even for firms with a full set of 156 observations, the model places some weight on modeled volatility. After all, even a sample of 156 observations contains sampling errors.

[15] EDF8 used three years of weekly returns to measure asset volatility in North America and Japan and five years of monthly returns elsewhere. This method was, in part, due to the differences in the quality of equity returns available in different parts of the world, historically. For more information, please refer to Knowledge Base FAQ438 (2009).

Modeled volatility is composed of three additive components, as we find three factors effective in explaining asset volatility: size, location, and business type.

As larger firms possess more diversified assets, they tend to have lower overall asset risk. Looking at a variety of subsets, we find that book value of assets, as well as sales, are two good indicators of size, but their relative importance in determining asset volatility varies across industries. For each industry and over a large panel of data, spanning multiple decades, we calibrate the non-linear relationship between asset volatility and firm size.

Each month, after calculating empirical volatility using the subset where at least approximately two-years of equity returns are available, we subtract the size-implied asset volatility from empirical volatility and cross-sectionally regress the residuals on country dummies and industry weights. Over time, the coefficients of these monthly regressions can be used to construct time series that capture the systematic asset risk for each country-industry pair. One can think of these as volatility fixed-effects.

After calculating the volatility fixed-effects and before calculating modeled volatility, we must account for small sample issues in specific countries. While, as of now, all 68 country groups are sufficiently populated, most have had periods of sparse data at some point during the past five decades. In such cases, we apply a Bayesian adjustment. We use five regions, and, for any country with less than 35 observations, after identifying its region and the fixed-effects of all countries in that region with more than 35 observations, we take the average of those fixed-effects as a Bayesian-prior and adjust the country's fixed-effect accordingly. For countries where no firms with at least two years of equity returns are available, the region level average fixed-effect is used until such firms become available.

As the description above shows, we calculate modeled volatility using two steps. We use a regression that estimates the country and industry fixed-effects after accounting for the size effect. The second step captures the systematic asset risk, which is cyclical. By cross-sectionally running the regression and updating the fixed-effects on a monthly basis, we can measure the state of systematic asset risk.

Modeled volatility does serve another purpose as well — facilitating the forward-looking adjustment. We elaborate on this point in the next section.

3.3.3 Forward-Looking Asset Volatility

What's New in EDF9: The forward-looking asset volatility adjustment is a new enhancement.

One of the completely new features introduced in EDF9 is forward-looking volatility. When aggregate asset risk in a region or industry has been far above or below its long-run level, we find it to mean-revert later. As the DD requires future asset risk, due to the backward-looking nature of any direct measure of asset volatility, we develop a forward-looking adjustment that results in a more accurate, yet stable, measure of credit risk. We model the necessary adjustment by looking at the gap between short-run and long-run measures of asset risk within a country and industry.

The rest of this section provides theoretical justifications and then proceeds with implementation details.

3.3.3.1 Theory

In its stylized form, we can split an asset return into systematic and idiosyncratic components.

$$\boldsymbol{r_{f,t} = \beta \cdot r_{s,t} + r_{i,t}} \tag{15}$$

Of the subscripts, f = "firm," s "systematic," i "idiosyncratic," and t "time." The coefficient β can be thought of in terms of a CAPM β, as it captures the sensitivity of the firm's returns to those of the market. Equation 16 implies a relationship among the standard-deviation of the different components:

$$\boldsymbol{\sigma_f^2 = \beta^2 \cdot \sigma_s^2 + \sigma_i^2} \tag{16}$$

As there is time-variation of asset risk, we introduce a time subscript:

$$\boldsymbol{\sigma_{f,t}^2 = \beta^2 \cdot \sigma_{s,t}^2 + \sigma_{i,t}^2} \tag{17}$$

We model the firm's idiosyncratic component as time invariant, while the systematic component is counter-cyclical, with mean-reverting tendencies.

In practice, our measurement of $\sigma_{f,t}^2$ at time t is based on the most recent three years of data, whereas, for modeling purposes, we are interested in asset risk in the future. Let us begin by assuming the existence of a long-run asset risk level for the asset return and its components at a time τ in the future.

$$\boldsymbol{\sigma_{f,LR}^2 = E\left[\sigma_{f,t+\tau}^2\right] = \beta^2 \cdot E\left[\sigma_{s,t+\tau}^2\right] + E\left[\sigma_{i,t+\tau}^2\right]} \tag{18}$$

Where the expectation is taken over τ (time). For shorter notation, we can write:

$$\boldsymbol{\sigma_{f,LR}^2 = \beta^2 \cdot \sigma_{s,LR}^2 + \sigma_{i,LR}^2} \tag{19}$$

If we assume that $\sigma_{i,LR}^2 = \sigma_{i,t}^2$, we can subtract equation 17 from equation 19 and obtain:

$$\boldsymbol{\sigma_{f,LR}^2 = \sigma_{f,t}^2 - \beta^2 \cdot \sigma_{s,t}^2 + \beta^2 \cdot \sigma_{s,LR}^2} \tag{20}$$

The relationship above shows that when current levels of systematic risk diverge from their long-run level, an additive adjustment is required to obtain the long-run level of a firm's asset risk. This approach is consistent with patterns found in the data. When volatility has been elevated over the past three years, it tends to be less elevated over the next year and vice-versa. Replacing $\boldsymbol{\beta^2}$ with $\rho^2 \cdot \frac{\sigma_{f,t}^2}{\sigma_{s,t}^2}$ we obtain:

$$\boldsymbol{\sigma_{f,LR}^2 = \sigma_{f,t}^2 - \rho^2 \cdot \frac{\sigma_{f,t}^2}{\sigma_{s,t}^2} \cdot \left(\sigma_{s,t}^2 - \sigma_{s,LR}^2\right)} \tag{21}$$

$$\sigma_{f,LR}^2 = \sigma_{f,t}^2 \cdot \left[1 - \rho^2 \cdot \left(1 - \frac{\sigma_{s,LR}^2}{\sigma_{s,t}^2}\right)\right]$$

$$\sigma_{f,LR}^2 = \sigma_{f,t}^2 \cdot \left[(1 - \rho^2) + \rho^2 \cdot \frac{\sigma_{s,LR}^2}{\sigma_{s,t}^2}\right]$$

To put into words, the forward-looking adjustment is a weighted-average of one and the ratio between long-run systematic asset volatility and short-run systematic asset volatility.

3.3.3.2 Practice

The final step before calibrating $\boldsymbol{\rho}$ is to choose statistical proxies for $\sigma_{f,t}$, $\sigma_{s,t}$, and $\sigma_{s,LR}$. Our measure for $\sigma_{f,t}$ is the weighted-average of empirical asset volatility and modeled asset volatility, regarded in EDF8 as asset volatility. For the adjustment ratio, we use modeled volatility to proxy $\sigma_{s,t}$ and long-run modeled volatility to proxy $\sigma_{s,LR}$.

Estimation of long-run modeled volatility is closely tied to the short-run modeled volatility (for brevity, modeled volatility). Modeled volatility is composed of three additive terms, a financial statement-implied volatility, estimated at the firm-level, a country-level, short-run, fixed-effect, and an industry-level, short-run, fixed-effect, weighted by the firm's percentage exposure to each industry. The estimation of short-run fixed-effects is updated once a month. For purposes of long-run modeled volatility, we take a 100-month moving average of the short-run, fixed-effect parameters, and we use the resulting fixed-effects to calculate long-run, modeled volatility. It is important for the long-run measure to not be conditioned on any particular phase of the business cycle, yet responsive to structural changes in an economy. Given National Bureau of Economic Research (NBER) data on recessions during the Great Moderation[16] suggest that recessions, peak-to-peak, have been about eight to nine years apart, we choose 100 months as the window size.

Last, we must calibrate $\boldsymbol{\rho}$. The constant $\boldsymbol{\rho}$ has two types of effects on the final EDF measures. First, it changes the rank-order of DD, and, second, it has a mild impact on the time series of aggregate EDF measures by way of lowering the amplitude of credit cycles. We find that, while not strongly identified, $\boldsymbol{\rho}$ is most suitably calibrated to 50%. Subsequently, this calibration results in a mild reduction of the amplitude of EDF cycles. In particular, among all the different time periods, we find aggregate EDF9 relative to aggregate EDF8 to be highest in the pre-financial crisis year 2007. This finding is primarily due to this adjustment, although not the only consequence of the forward-looking adjustments.

In fact, adjusting volatility to be more relevant in the future has increased the model's early warning capabilities. In other words, EDF values of high-risk firms begin to rise earlier, hence, separating themselves from non-defaulters far before default.

[16] In this paper, the Great Moderations refers to the economic era started in the mid-1980s where both the economic cycles and inflation levels moderated.

Figure 2 shows the effect of this adjustment on aggregate EDF levels. The green line is not subject to any adjustment, and the blue line shows the adjustment as calibrated under EDF9. It is noteworthy that the adjustment ratio raises EDF levels during the pre-crisis year of 2006, while lowering it in the post-crisis year of 2010. As the population sample used in aggregation does change over time, it is helpful to look at a firm-level example. Figure 3 shows how different levels of forward-looking adjustment affects a firm's EDF measure.

Figure 2 Effect of Forward-Looking Adjustment on Aggregate EDF Time series

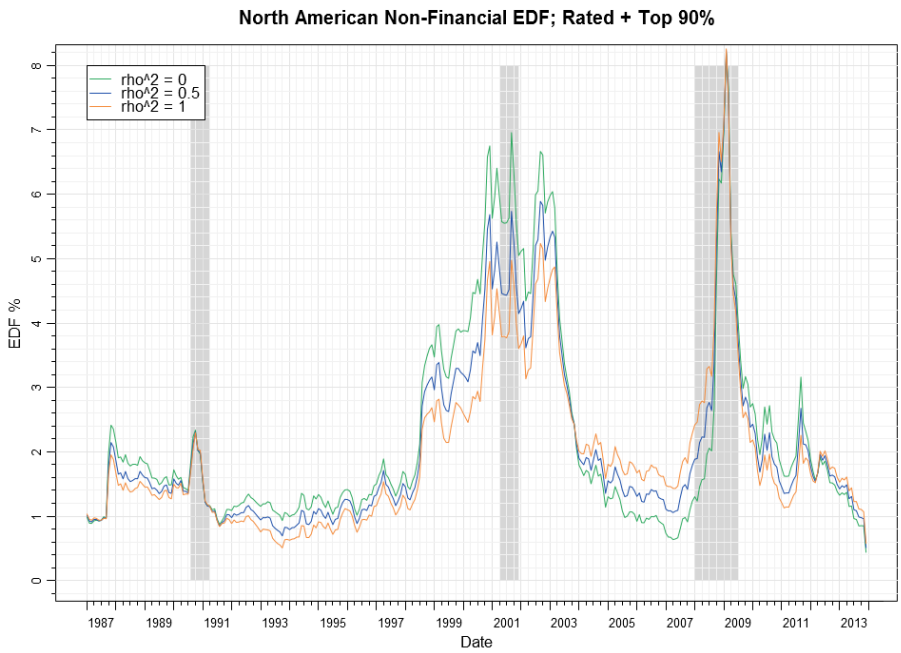

Figure 3 Effect of Forward-Looking Adjustment on Firm-Level EDF Time series

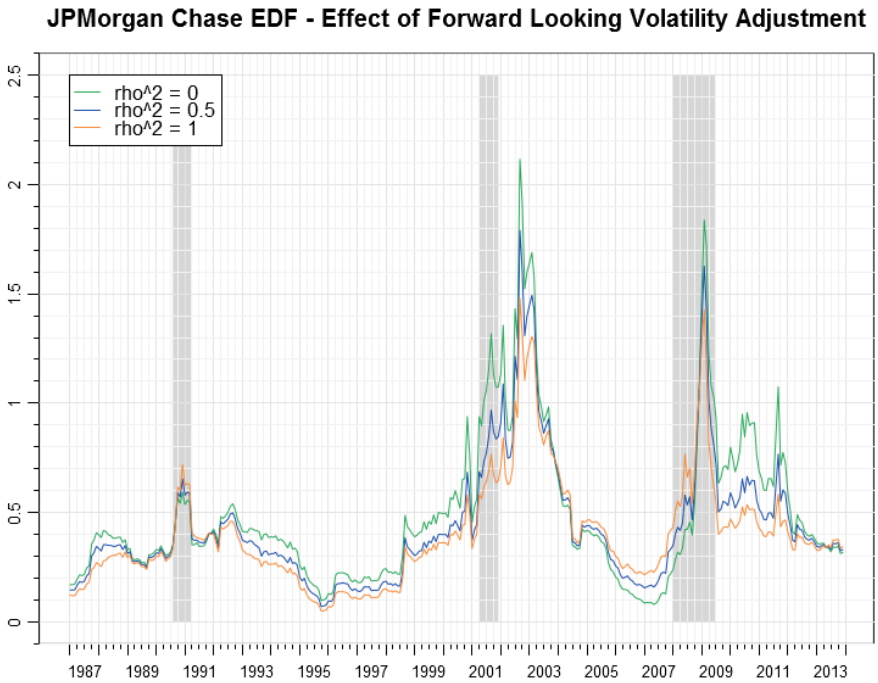

4. Mapping Distance-to-Default to EDF Measures

This section explains how the EDF9 model transforms the Distance-to-Default to produce a probability, a key output. Since this transformation is the most important calibration of the model, in the next subsections we cover the DD's theoretical foundation, conceptual grounding, and implementation considerations. With EDF9, we update the calibration, incorporating data from the most recent credit cycle. Further, we introduce a new financial firm specific mapping. Finally, we increase the upper bound of the EDF to 50% for non-financial firms, and we discuss the reasoning behind these changes

4.1 Concept

The Distance-to-Default concept provides an ordinal measure of default risk. In theory, the concept has the interpretation of a standard deviation, so one could infer a probability. If the underlying stochastic process had a Gaussian distribution, then we can estimate the one-year EDF measure as:

$$\Phi\left[-\frac{ln\left(\frac{A_0}{D}\right)+\left(\mu_A-\frac{1}{2}\sigma_A^2\right)}{\sigma_A}\right] \tag{22}$$

In practice, we obtain more accurate results by empirically transforming the DD to the physical probability space, with a monotonic function *M*()

$$\boldsymbol{M}\left[-\frac{ln\left(\frac{A_0}{D}\right)+\left(\mu_A-\frac{1}{2}\sigma_A^2\right)}{\sigma_A}\right] \tag{23}$$

We call the result the EDF credit metric.[17]

In the formula above, strictly speaking $\frac{ln\left(\frac{A_0}{D}\right)}{\sigma_A}$ is the asset's distance from default as of today. That distance is projected to become $\frac{ln\left(\frac{A_0}{D}\right)+\left(\mu_A-\frac{1}{2}\sigma_A^2\right)}{\sigma_A}$ by next year. When we say DD, what we intend is this projected value. Uncertainty surrounds this projected value. The uncertainty is due to the economic shocks, both systematic and idiosyncratic, that affect the firm. Theoretically, the cumulative distribution function for that uncertainty is the cumulative standard normal distribution

The DD-to-EDF mapping $M(\,)$ is our calibration for the distribution of these economic shocks and is perhaps the primary calibration of the model.

4.2 Calibration

When evaluating a model, an important question to ask is whether the calibration is relevant to "my" portfolio in the future. There are two aspects to this question: a cross-sectional one and a temporal one. "My portfolio" is a certain cross-section of all firms. Temporally, the calibration is to be used "in the future," while it was estimated against historic data.

A portfolio can be characterized by its geographic breadth, as well as the business types it covers. For instance, one may have a portfolio of firms concentrated in Southeast Asia, while another may have a portfolio of firms in Western Europe and South Africa. As for business type, one may have a portfolio of only non-financial firms, whereas, another portfolio may be concentrated in banks. Further, there is the idiosyncratic risk of individual names within the portfolio. Is $M(\,)$ relevant to different cross-sections?

The answer is in the conceptual role of $M(\,)$. The Public Firm model produces probabilities, but with only a single factor. Hence, calibration only involves estimating the distribution of the shocks. More specifically, we calibrate the model to parameterize the unanticipated component of default risk, while the factors in DD capture the anticipated and measurable ones. The anticipated ones have become anticipated because a shock has been realized or a policy has been made and, hence, are observable.

For instance, after economic troubles hit Southeast Asia in 1997, investors anticipated increased risk, reflected in lower stock market valuations. Such valuations reduce A_0, one of the factors in DD. But in 1996, Southeast Asia troubles were not anticipated, just as in early 2008, when the Great Recession was not anticipated.[18]

On the one hand, a bank with a concentrated portfolio will possess risky assets that lead to a relatively high estimation of asset volatility, which, in turn, is reflected in DD through σ_A. On the other hand, when a firm in need of cash taps into its credit line, we

[17] Crosbie and Bohn (2003) discusses the unsuitability of a Gaussian distribution.

[18] On August 29, 2008, the market did not anticipate a 30% reduction in the S&P 500 Index over the following three months.

can observe its increased liability, reflected via the default point. In either case, DD captures the observable risk, while $M(\,)$ reflects the risk of future shocks.

We may expect oil and gas exploration firms to experience elevated credit risk when energy prices fall unexpectedly. We may expect the future earnings potential of such firms to decline with energy prices. To the extent that the market expects a decline in future earnings, the market capitalization of such firms would decline, which would lead to a decline in DD. The calibration of $M(\,)$, in turn, captures the likelihood of asset values falling in the future, with sufficient severity to lead to default.

Unanticipated shocks need not be systematic and can hit individual names as well. For example, the EDF value of a specific firm may jump when a lawsuit is announced, which leads to a decline in stock price. We can interpret the calibration of $M(\,)$ as the parameterization of the distribution of both systematic and idiosyncratic credit risk shocks. In other words, $M(\,)$ captures the magnitude of different shocks and the frequency at which they hit. As shocks are not bound by geography or business models, we assess the calibration to be relevant across macroeconomic environments.

In making a calibration relevant for use in the future, in our view, it is better to calibrate to a longer history.

If we calibrate to only one credit cycle, we would see many idiosyncratic firm-level shocks, while only having a limited number of aggregate shocks. As we include more history, we include more systematic shocks into the calibration, which provides a more robust calibration compared to a model that uses a relatively shorter window. By extension, if a certain credit cycle is included, it is important to cover it in-full, as partial inclusion can skew the observed distribution of shocks.

A firm's DD is intended to control for the future earnings potential of the company as anticipated by the equity market. Moreover, the DD reflects the firm's business practices, the industry and the country that the firm resides in, and the current state of the economy. What DD cannot control for are the unanticipated economic shocks — both aggregate and idiosyncratic. The DD-to-EDF mapping intends to capture the magnitude and likelihood of these unanticipated events. In our view, a larger calibration sample makes for a more robust DD-to-EDF.

4.3 Implementation

What's New in EDF9: *We update the DD-to-EDF calibration by adding the most recent credit cycle*

The process of building the DD-to-EDF mapping involves determining the frequency of realized defaults, conditional on the credit risk level. Figure 4 shows the EDF credit metrics and observed default frequency (ODF) for 50 different credit risk levels, determined by grouping 25 years of rated, non-financial corporate observations into equally-sized buckets. Figure 4's blue line can be thought of as the DD-to-EDF mapping, had DD's been discretized at 50 levels. The orange dot represents the percentage of defaults within each bucket, and, like the estimator for the binomial parameter, its standard error, in part, depends on the number of flagged observations. Hence, we see more sampling variability on the right end of the spectrum, when compared to the left.

The blue line is a flexible function form with sufficiently many degrees of freedom to fit through the orange dots. The parameters are estimated, such that, the mapping function performs well on a variety of different subsets.

At any point on the credit risk spectrum, the mapping function intends to estimate a probability of default that is consistent with the realization of default over an extended period of time.

Figure 4 DD-to-EDF Mapping vs. Observed Default Frequency

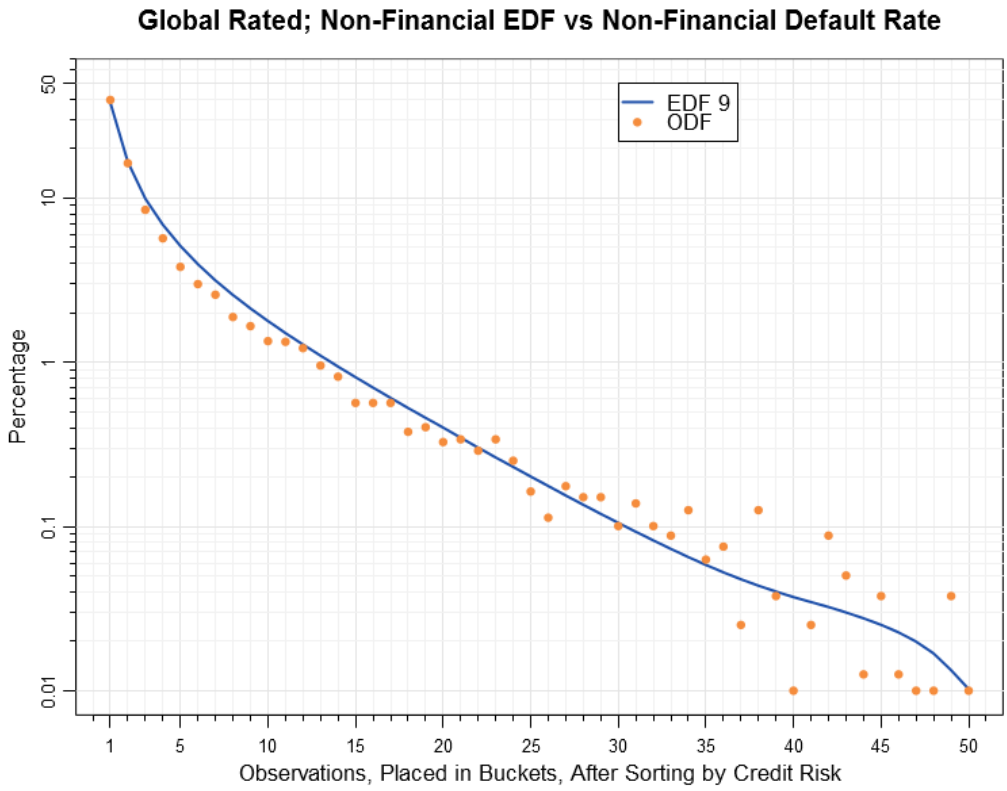

4.4 Defining Default

We regard a company in default if it does not make a payment on a debt obligation when due at maturity date. We record the event as the date of the missed payment and do not take into account grace periods given by the lender. This definition is more conservative relative to Basel II, which regards 90-days past due as default.[19] When a lender extends the maturity date of a debt while the company has limited access to funds and limited profit generation, if we find that to have been to prevent bankruptcy or default, we regard it as a default. A distressed exchange is also treated as a default, where debt is exchanged for equity and a loss of seniority or a decrease in principal or interest amount occurs. Furthermore, the default of a wholly-owned subsidiary that represents more than half of the sales or assets of the public parent company is regarded as default of the parent company. Last, we keep track of bankruptcies and liquidations as well.

For financial firms, we consider government bailout as a type of default. We look for indications as to whether the firm would have defaulted without the received support. Many banks receive government funding for a variety of reasons. We only consider the receipt of government funds as a bailout if, in our judgment, the funds are provided for the purpose of avoiding a default on the part of the bank.

4.5 Financial Firms

What's New in EDF9: *We introduce a financial firm-specific DD-to-EDF calibration.*

In Section 3.2.1, we discussed how our approach results in a different default point definition for financial firms. Here, we discuss the implications to the DD-to-EDF measure mapping.

A financial firm's financial statements are fundamentally more opaque than non-financial firm statements. This opaqueness affects the DD's information content and, subsequently, its conditional distribution. For the purpose of DD-to-EDF measure mapping, this difference implies that a calibration sample including non-financial corporations should not be used. If we preserve the relation that Figure 4 implies between DD and EDF measure, but change the sample to financial firms, we obtain Figure 5

[19] Refer to Basel Committee on Banking Supervision (2004).

Figure 5 DD-to-EDF; Financials vs. Non-Financials

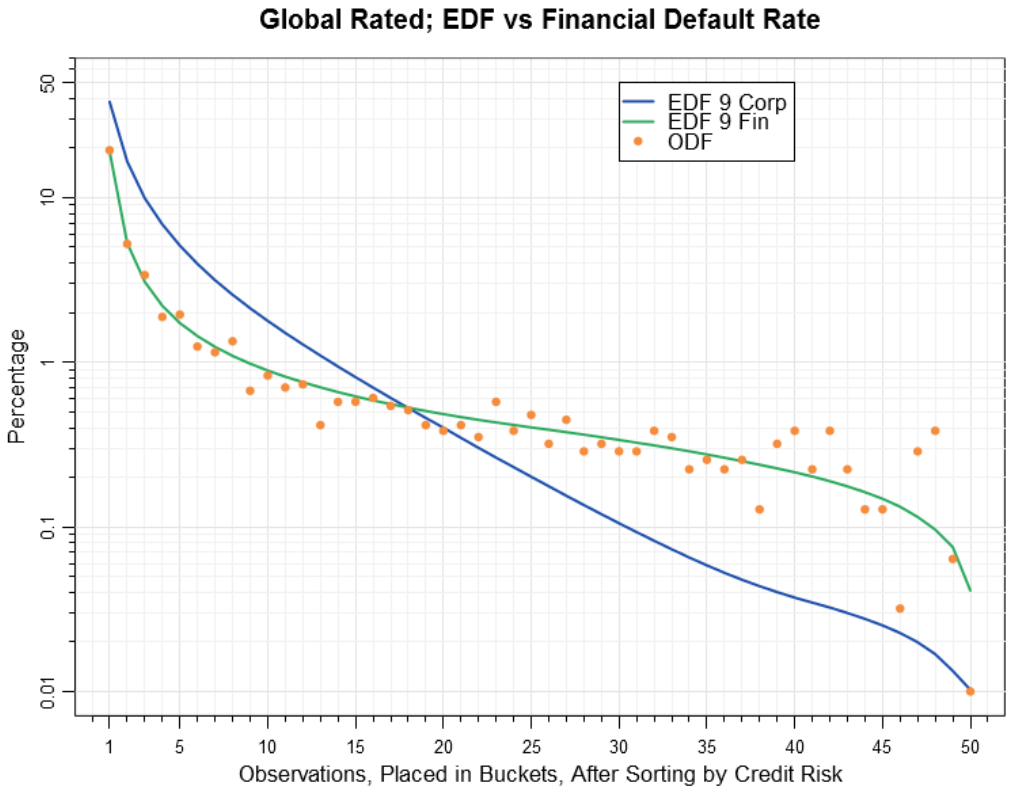

Figure 5 is significant for many reasons. First, while there are cases of default among publicly traded financial firms in the past, prior to the financial crisis in North America and Europe's double-dipping financial crisis, there were not enough observations to identify the relationship that the orange dots in Figure 5 demonstrate.

The second significance is the implied slope. As the opaqueness of a financial firm's statements reduces the signal-to-noise ratio of the market signals, the EDF measure would have to be less sensitive to market changes than in the case of non-financial corporations. The default data confirms that a change that moves a financial firm's DD from, say, bucket 30 to bucket 25, warrants a smaller increase in the EDF measure than it would have had it been a non-financial firm.

As a benefit of having a financial firm-specific mapping, we do see more stable EDF measures in EDF9 when compared to EDF8. This finding implies that the riskiest financial firms in a bad stage of the cycle look safer under EDF9, but it also implies that the safest firms in a good stage of the cycle look riskier under EDF9. This result follows also, in part, because the sample includes bailouts, which can raise default frequency for the low-risk DDs. We include bailouts in the sample, as the model reflects the firm's stand-alone credit risk. Over time, as macroeconomic fluctuations move the distribution of firms across the DD spectrum, EDF9 may result in a more stable measure of portfolio level risk and capital requirements for financial firms.

4.6 Upper Bound

What's New in EDF9: *We increase the upper bound to 50% for non-financial corporations, but keep the upper bound at 35% for financial firms.*

EDF9 provides extra granularity for high-risk corporate names when compared to prior versions. Firms in financial distress may avoid default in several ways. One common way occurs when another firm buys the distressed firm at a low price, assumes ownership of the assets, and honors the debt (e.g., JPMorgan's acquisition of Bear Stearns or Wells Fargo's acquisition of Wachovia). Consequently, prior to the default event, the EDF measure is always less than 100%, and the upper bound on the EDF metric is an empirical question. By studying graphs similar to Figure 4 but with more granularity in the number of buckets, we

conclude that capping the EDF metric at 35% results in an under-estimation bias for EDF9. In other words, the model's separation power among the very high-risk names has increased so much, it is appropriate to distinguish between observations with a 35% credit risk and those with a 50% credit risk.

Figure 6 DD-to-EDF Measure; Upper-Bound for Non-financial Corporations

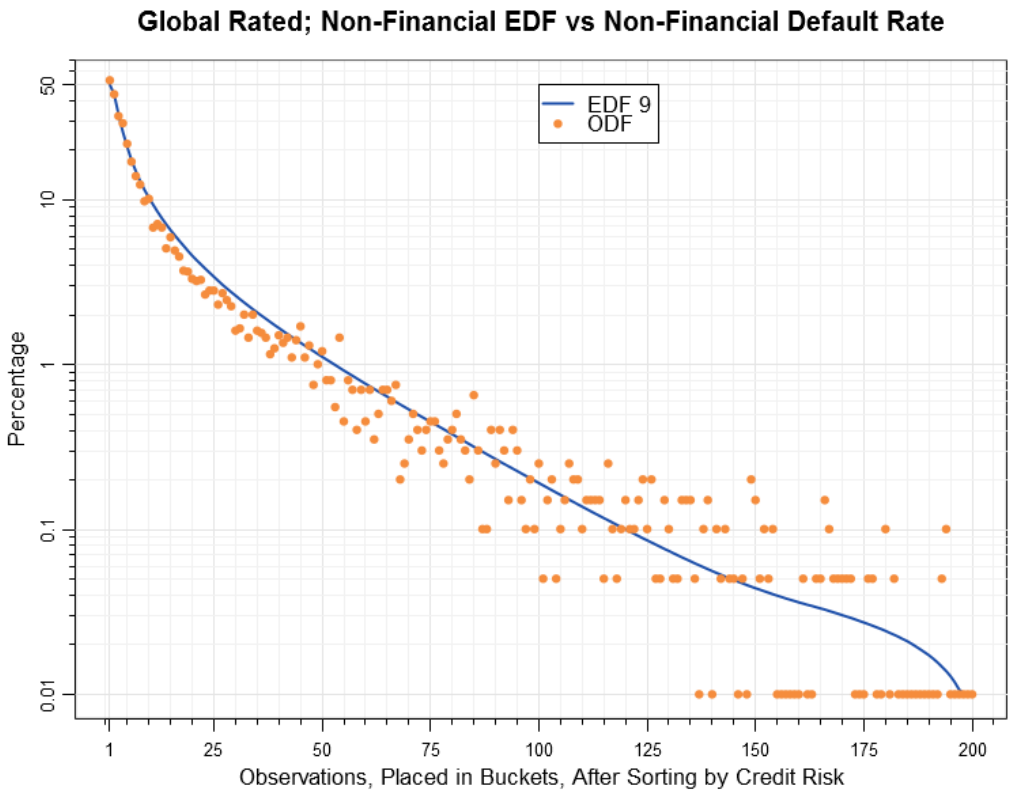

Figure 7 DD-to-EDF Measures; Upper-Bound for Financial Firms

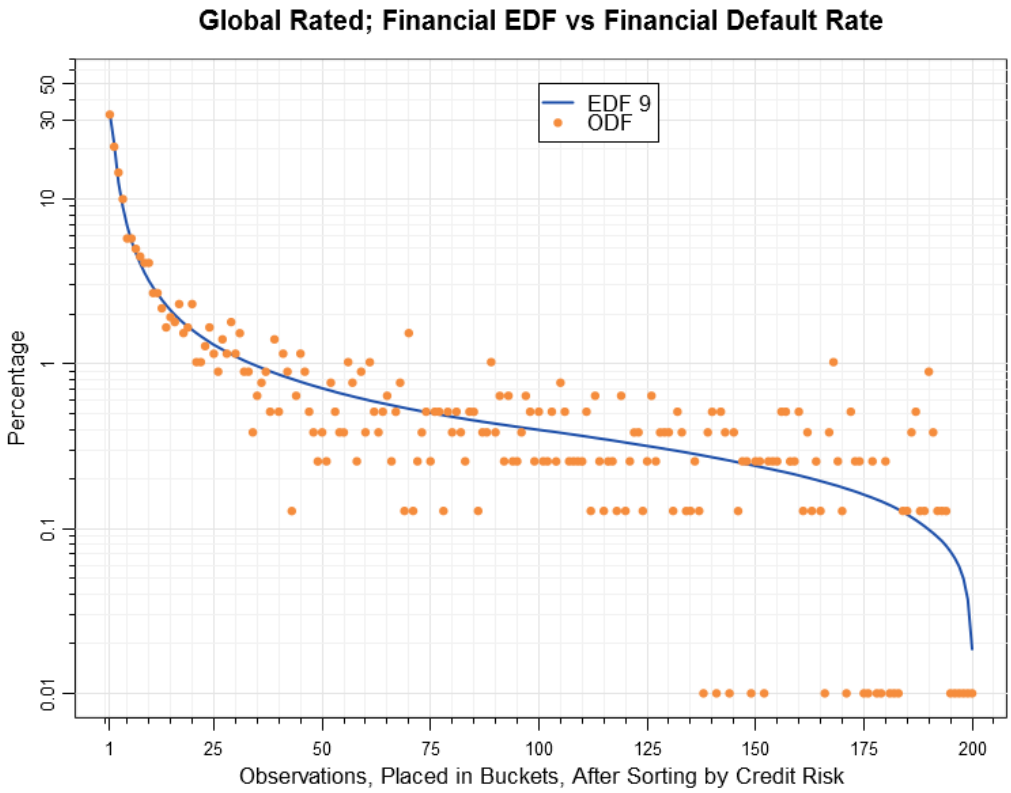

Figure 6 shows what Figure 4 would have looked like, had we used 200 buckets instead of 50. Each DD observations bucket is now 1/4th smaller and sampling variability is naturally more pronounced. But focusing on the left end of the graph, the first orange dot of Figure 4 has now split into the first four orange dots of Figure 6. Graphs similar to this one, created with a variety of different bucketing schemes, allow us to study the EDF level for the highest-risk names. The EDF measure is now capped at 50% for non-financial corporates. Figure 7 shows the same study for financial names. In this case, we can see that, consistent with the flatter mapping function, the EDF level does not peak as high as it does for non-financials. As suggested by the data, the cap remains at 35% for financial firms.

4.7 Performance in Different Portfolios

On the one hand, many models in the option pricing literature — the Merton model for instance — have made the Gaussian assumption for this distribution. On the other hand, the data consistently rejects this assumption. Instead of the cumulative normal distribution, the Public Firm EDF model uses physical realization of default to obtain the most accurate assessment of the relation between the DD and EDF measure. This process requires the collection and construction of a database of default events, an endeavor that Moody's Analytics has been engaged in for more than two decades. This research creates a database of 10,000-plus default events, temporally spanning many decades and cross-sectionally, covering firms across the globe.

Default collection is a challenging and resource-intensive endeavor, for many reasons. There are more than 35,000 firms, spread globally, and the breadth of the sample is rapidly expanding. Second, firms are subject to different disclosure requirements and operate under different legal environments. The information is also disseminated in a variety of different languages. Furthermore, operationalizing default definitions for small firms, as well as those subject to government intervention, can be challenging. All of these issues impose notable restrictions on the calibration sample of DD-to-EDF measure mapping. As important as temporal depth and cross-sectional breadth of the sample are, it is also important to make sure default collection is comprehensive. We restrict the calibration sample to a subset of the overall population to ensure that the default collection is reliable.[20]

[20] For more on challenges of default collection, please refer to Section 3.1.2 Dwyer and Qu (2007).

In the universe of rated firms, we view the coverage of defaults as comprehensive. Among non-rated firms in North-America and Europe, we view the coverage to be close to comprehensive for firms among the top-90% of liabilities. In Appendix C, we discuss how we define the top-90% and what it means for firms of different sizes.

For any sample chosen for such calibration, sampling variability is always present. To properly account for this issue, we construct many different samples with sufficient depth and breadth. We then ensure that, for any point on the credit risk spectrum, the EDF measure is an appropriate reflection of credit risk, allowing for sampling variability.

Figure 8 shows some of this effort. The left panels shows sampling variability resulting from looking at different recessions, within different regions. The right panels show looking at different cuts of firms, within different regions. Specifically, the right panels show the DD-to-EDF measure mapping properly captures the slope and curvature for the different samples. We find the different levels suggested in the bottom right graph reflects the challenges involved in comprehensive default collection for that sample, but it also indicates that the slope and curvature of the mapping function is just as prudent for the top-99% sample as it is for the top-90% sample.

EDF9, when compared to its predecessor, benefits from more sample depth and breadth. The EDF9 data includes the Great Recession, as well as its recovery years. Furthermore, for the first time, the Public Firm EDF model expands the calibration sample to include global firms. We now include any rated firm, regardless of location or size, in the calibration sample.

Figure 8 Sampling Variability in Observed Defaults

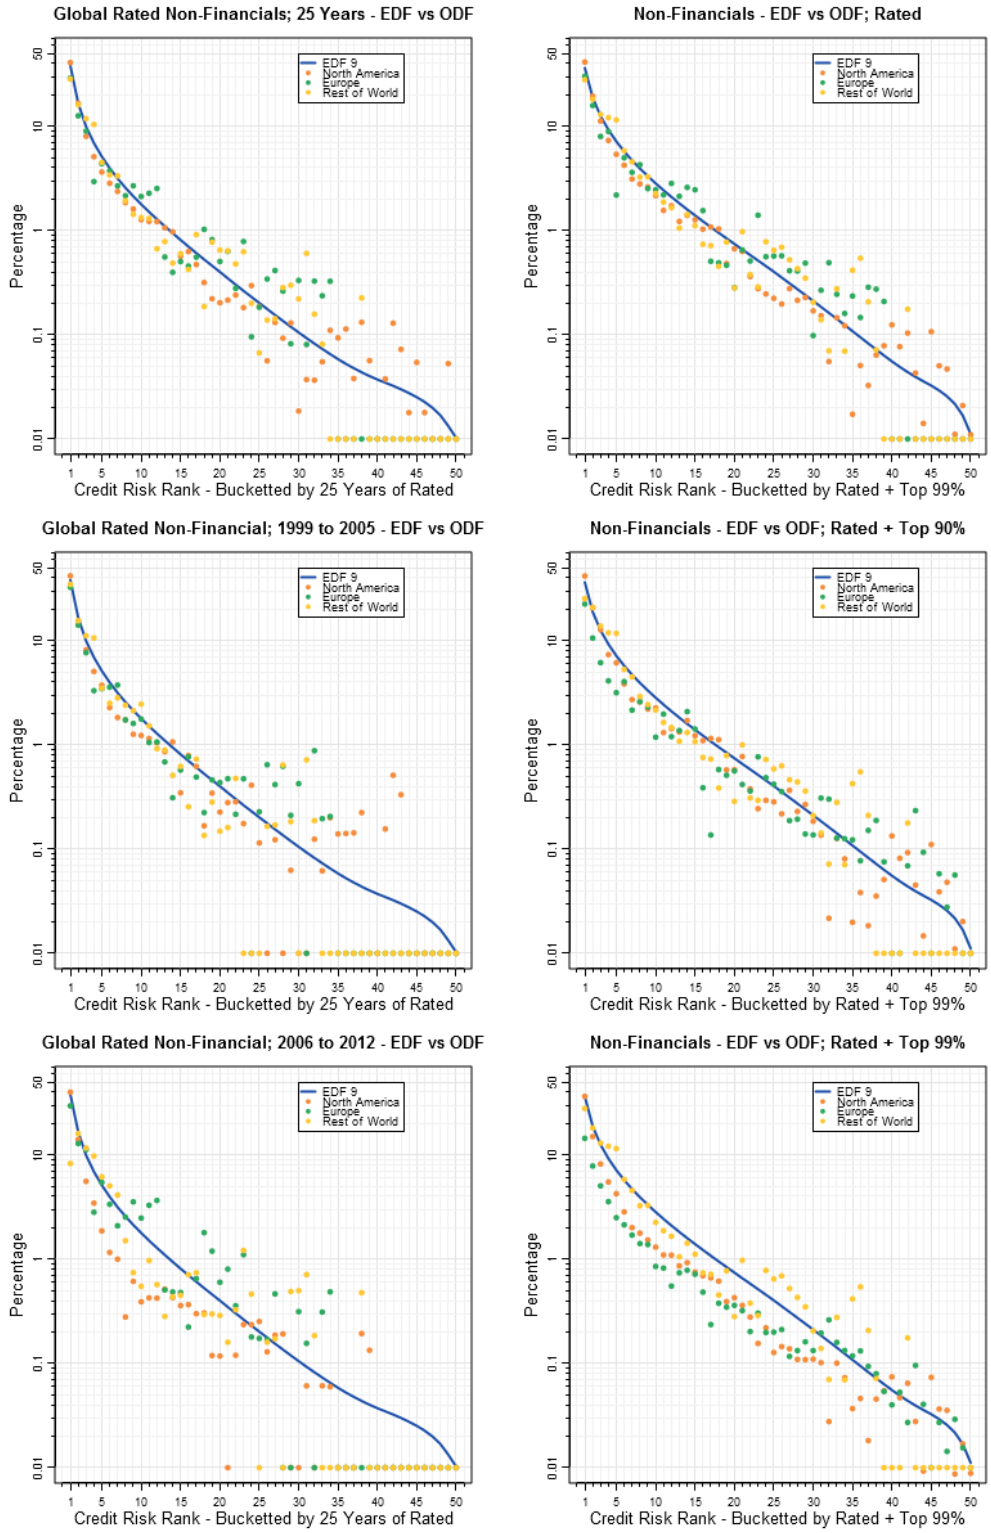

5. Assessing Public Firm EDF Model Performance

There are many performance metrics with which a credit risk model can be assessed. This section presents the conceptual grounding for some of the tests we perform and shares some results. For more details, please see Moody's Analytics "Public EDF Validation" (2015).

Assessing the performance of the ordinal measure of default risk is a different problem from that of the cardinal measure of default risk. In other words, producing prudent probability levels does not imply high rank-ordering power, and inversely high rank-ordering power does not imply properly calibrated PD levels. As a result, we have a set of tests for Distance-to-Default assessing the model's rank-ordering power and another set for the EDF measure, assessing the probability levels.

5.1 Rank-Ordering Power of Distance-to-Default

The Accuracy Ratio (AR) measure produces a ratio between -100% and 100%, quantifying a model's accuracy in separating default observations from others. An AR of 100% is a "perfect model," and a model no better than a random model has an AR of 0%. A negative AR indicates that the model is calling safe firms risky, and vice-versa.[21]

Take for example a model rank-ordering twelve hypothetical firms.

Table 1

Rank-ordering of 12 Hypothetical Firms

FIRM	A	B	C	D	E	F	G	H	I	J	K	L
Defaulte	No	No	No	Yes	No	Yes	No	Yes	No	Yes	Yes	Yes
PD	1bp	3bps	10bps	40bps	70bps	1%	2%	5%	10%	20%	30%	50%

Further, suppose we have a probability of default model that enables us to rank-order the firms by the produced PD. There are more defaulted firms among the high-risk names and fewer among the low-risk names, so the model is performing as expected, but it is not perfect. The AR helps us quantify this fact.

Of the twelve names, six have defaulted, so ideally, when looking at the riskiest half of the subset, we should find all the defaulted names. And when looking at the riskiest quarter, we should find half of the defaulted names. We can plot the percentage of observed defaults below each risk level, for a variety of different risk levels.

Ideally, a model rank-orders perfectly and produces results such as the green bars shown in Figure 9. Realistically, the noise and uncertainty in the future result in model performance that does not perfectly overlap with the green bars. A random model results in a lower triangular shape, with the top of its bars connecting the lower left corner to the upper right corner with an approximately straight line.

[21] Traditionally, the Accuracy Ratio is derived from a Cumulative Accuracy Profile (CAP). Closely related to a Cumulative Accuracy Profile is a Received Operating Characteristics (ROC) curve. Typically, a Gini coefficient is derived from an ROC curve. A Gini coefficient can be shown to be equal to an AR.

Figure 9 Cap Plot for the Hypothetical Example

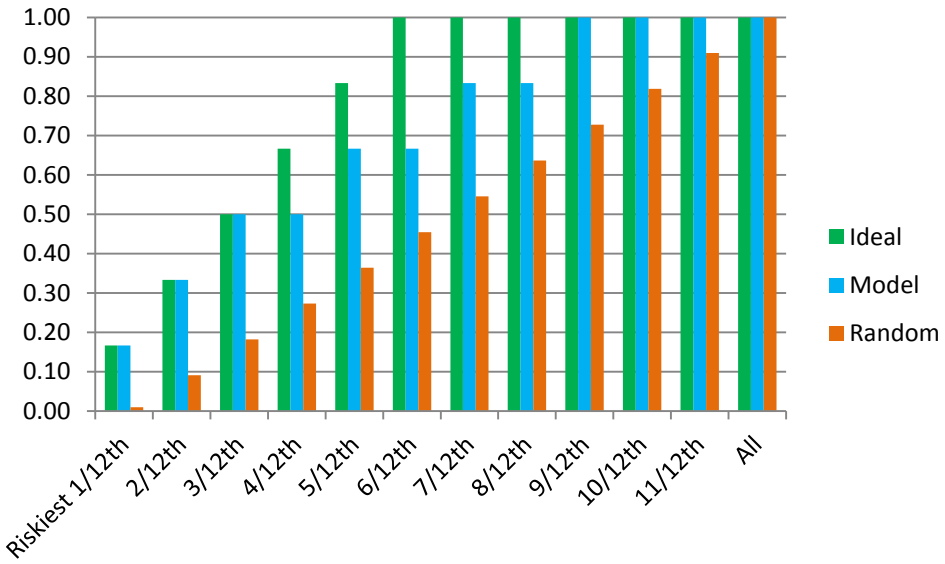

By convention, the ideal model has an accuracy ratio of 100%, and the random model has an accuracy ratio of 0%. For anything in between, the value of the accuracy ratio is the percentage of overlap between the model's bars and ideal bars, for the area bound between the two extremes of ideal and random. In our case, the blue bars have a 67% overlap with the area bound between the top of red bars and the top of green bars. As seen in Figure 9, the bars form the Cumulative Accuracy Profile, also called the CAP plot.

For a real-life case, Figure 10 looks at a set of global rated firms, spanning 1987 – 2012, resulting in about 600,000 observations, about 1.6% of which are default observations.

Figure 10 Global Rated Firms, 1987 – 2012

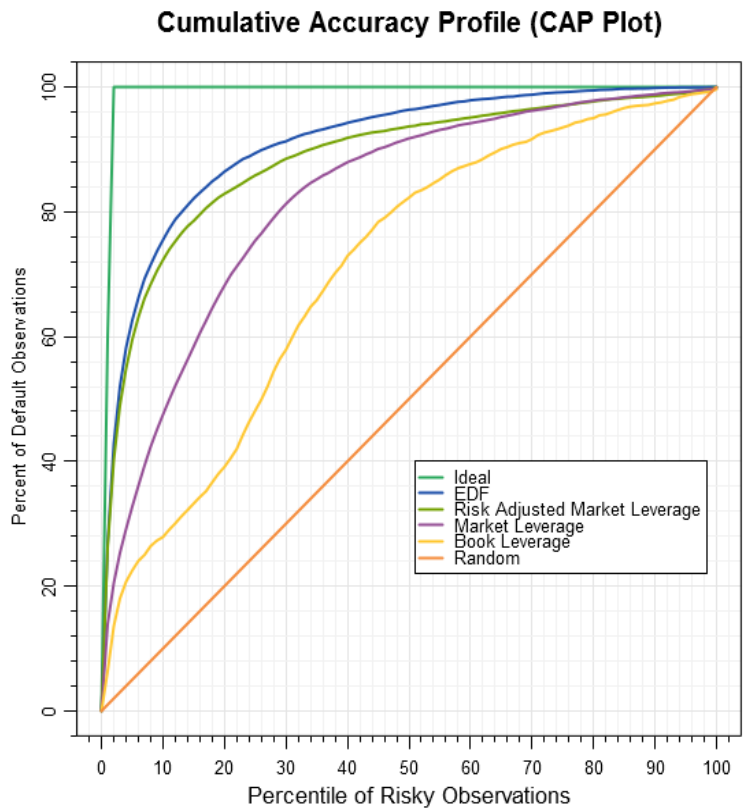

Next, we rank-order the firms using four different statistics. The first is the book-equity / book assets ratio introduced in Section 3. In the plot, it is called "Book Leverage." The second measure is market capitalization as a percentage of firm value, coming from replacing book values of equity with market values of equity. This statistic is labeled "Market Leverage" and shows that a market-based measure of equity is more useful for understanding credit risk than book equity.

The third statistic is the risk-adjusted measure of market leverage, defined as:

$$\frac{A - D}{A \cdot \sigma_A}$$

In Appendix B, we show that this ratio is an approximation of DD, and it can be thought of as a risk-adjusted measure of market leverage. The numerator is the difference between asset value and default point, and the denominator is asset value multiplied by asset volatility. The CAP plot shows it performs better than market leverage. The AR of this measure is the area between the dark green line and red line, which, against this dataset, happens to be 78%.

For the fourth statistic, we look at the EDF measure. We see that the EDF curve is closest to the ideal line and furthest away from the random line. The accuracy ratio of the EDF measure is the area between the blue and red line, as a percentage of the area between the green and red lines, and, for this particular dataset, equals 83%.

5.1.1 Important Accuracy Ratio Caveats

It is important to note a few caveats regarding AR. First, when comparing the AR of two different models, it is important to ensure that they are measured on the same sample, as the particular value of AR depends on both the model and the sample. If we apply the same model to two samples, where the one sample has firms with very similar levels of risk, and the other has very different levels of risk, the AR of the sample with very different levels should be higher.

Second, AR does not make a statement about the model's out-of-sample properties. The out-of-sample properties must be studied separately. In other words, a model can achieve a very high AR on a specific dataset, but perform poorly on other, different

datasets. For a structural model, the out-of-sample rank-ordering properties tend to be strong, as there are very few calibrated parameters. The out-of-sample performance is more of an issue when comparing a structural model to an econometric model or when comparing a structural model to a model that uses machine learning techniques.

Third, AR is a statistic with a standard error and should be interpreted with regard to that. For example, an EDF measure's AR of 83%, while larger than risk-adjusted market-leverage's AR of 78%, may not be statistically significantly different.

Fourth, AR can be studied at a variety of different horizons. While we can study a model's separation power upon default, we can also do so six months before default or 36 months before default. What we learn from these studies is the model's early warning capability. Typically, a model has a higher AR when predicting default at shorter horizons.

5.1.2 Importance of Accuracy Ratio Standard Error

In this section, we stress the importance of understanding the distribution of accuracy ratio. For more technical details on the accuracy ratio, its statistical cousins, as well as distributional characteristics, refer to Engelmann (2003).

Table 2

Hypothetical Example on Rank-Ordering

FIRM	A	B	C	D	E	F	G	H	I	J	K	L	AR
Defaulted	No	No	No	Yes	No	Yes	No	Yes	No	Yes	Yes	Yes	
PD	1bp	3bps	10bps	40bps	70bps	1%	2%	5%	10%	20%	30%	50%	
Subset 1	x	x	x	x	x	x	x	x	x	x	x	x	67%
Subset 2			x	x	x	x	x	x	x	x			25%
Subset 3	x	x	x		x			x		x	x	x	100%
Subset 4					x	x	x	x	x				-33%

Let us begin with our earlier example. While the AR of our hypothetical PD model was 67%, that value differs against different subsets of the data.

In Table 2, four different subsets of the data result in four very different AR values, even one producing a negative value, suggesting the rank-order is incorrect and worse than a random order. Table 2 illustrates the effect of sampling variability on AR.

Figure 11 Bootstrapped AR (left panel); Early-Warning at Longer Horizons (right panel)

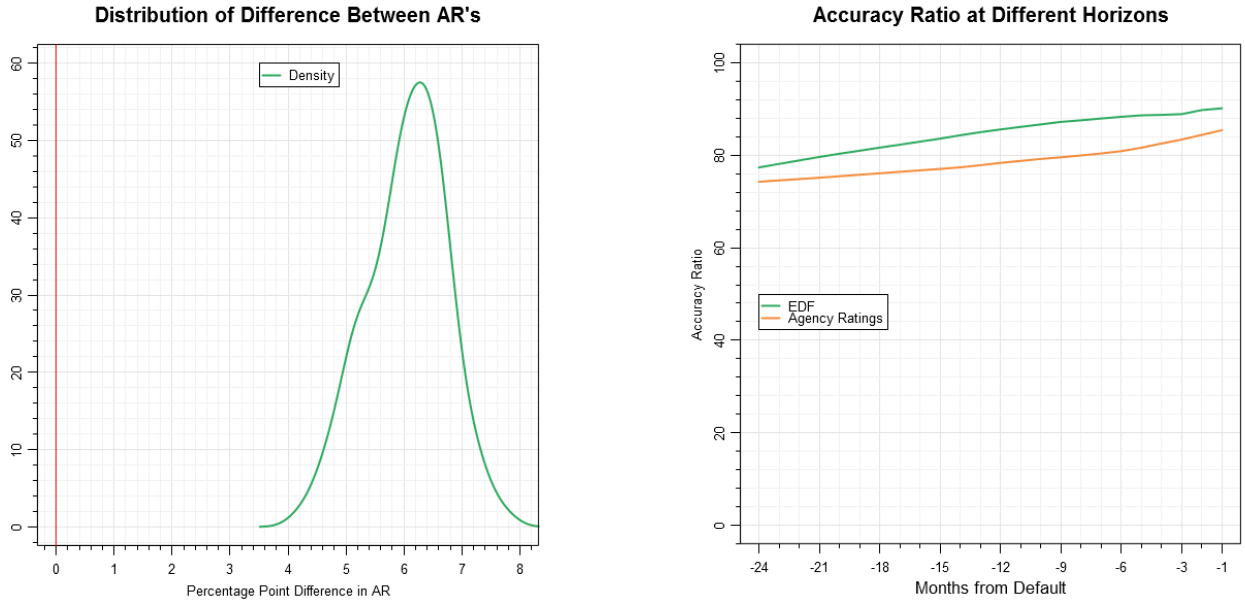

One solution to this issue is to bootstrap the difference between the ARs of different samples. Take for example rated corporations during the recent business cycle. If we take the more conservative of Moody's and S&P's ratings, we obtain an AR of 78%. If we take the EDF measure, however, we obtain an AR of 86%. But are these measures statistically significantly different? We can bootstrap the distribution by constructing, say, 100 different random subsets. On each subset, we calculate the AR using both rating and EDF measure and note the difference. We then identify the percentages that were positive to determine significance. In Figure 11 (left panel), we fit a non-parametric distribution to the observations. This figure illustrates the test, showing the null-hypothesis was rejected.

If we obtain a distribution centered around zero, it means that, even though the two methods have different rank-ordering powers on different subsets, one is as likely to be better as it is to be worse. Hence, neither model can be declared better that the other.

In the case of EDF measures vs. agency ratings applied to the non-financials during the recent business cycle, we see in Figure 11 (left panel) that the distribution is significantly to the right of zero, indicating we can infer EDF measures have a higher rank-ordering power.

5.1.3 Accuracy Ratio at Different Horizons

While we calculated all the AR's reported earlier in this section over a 12-month horizon, there are valuable details in studying other horizons as well. Figure 11 (right panel) presents the accuracy ratio as a function of the months from default. We calculate the ARs as follows. We construct a sample in which we include all the good observations, but for each defaulting firm we only retain the observations within X months from the month of the default event. In Figure 11 (right panel), the X axis represents the number of months between the observation date, and the y-axis is the corresponding accuracy ratio. We see one-year ARs of the previous section as y-values of 86% and 78%, corresponding to where the x-axis is labeled -12. As we move further away from default, we see a reduction in the AR.

When analyzing a model with such graphs, there are two aspects to note: the level of the line and its slope. A higher level indicates higher rank-ordering power, controlling for horizon. The slope indicates the pace at which early-warning power is lost. When comparing two models, it is useful to note whether their lines cross and, if so, at what horizon. When the lines of two models cross, it indicates one is better for early-warnings, while the other outperforms for firms closer to default.

5.2 Temporal Characteristics of EDF Levels

As discussed in Section 4, by construction, the mapping function is consistent with physical default rates over an extended period of time using a panel dataset. One can also compare the time series average of EDF measures to the observed default frequencies (ODF).

While an intuitive concept, defining ODF concretely is more challenging than may initially appear. "Default Intensity" is one way to define ODF. Before trying to interpret the graph, it is important to understand this definition and its limitations. Figure 12 compares one definition of ODF against EDF. To construct the orange line, each month, we look at the percentage of firms with an EDF at the end of the month and compute the percentage that default in the next month. Then we fit a Hodrick-Prescott (HP) filter through it to get the default-intensity.

Figure 12 Time Series of EDF vs. Default Intensity

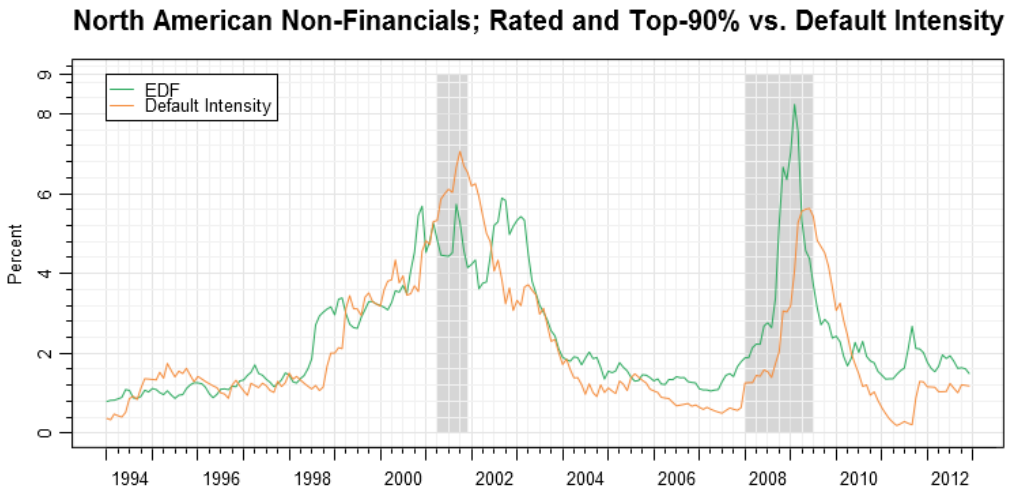

We use an HP filter because the default data is otherwise quite noisy. Figure 13 shows the "Instantaneous Default," which is the annualized, monthly observed default rate in Figure 12, but before applying the filter.

Figure 13 Time Series of EDF vs. Instantaneous Defaults

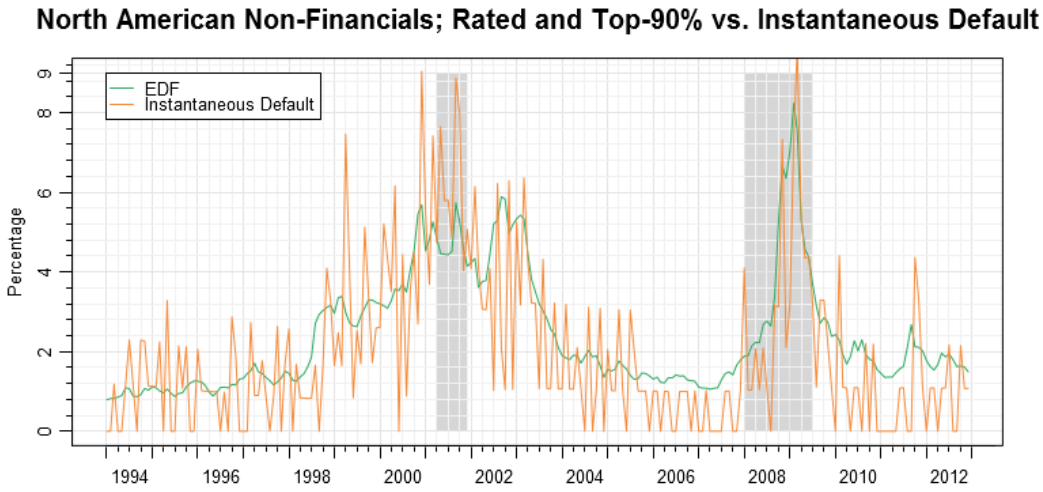

The orange lines in Figures 12 and 13 represent the same underlying information, however one masks the sampling variability in the data by smoothening the time series. In interpreting Figure 12, there are two aspects of the model to note: its lead against defaults and its level compared to defaults. There are three notable periods of sudden increase in risk, 1998, late 2000, and 2008. Figure 12 shows the model leading the default events by more than a quarter. To interpret the information regarding level, we must be cognizant of three things.

First, while Figure 12 suggests the EDF levels are appropriate in the left half but high in the right half, we must not forget that the sampling variability presented in Figure 13 is actually hidden in Figure 12. Sampling variability causes noisy default data. The model, however, intends to capture the underlying default intensity, and, by visual inspection of Figure 13, the model is fairly successful at that. To appreciate the importance of sampling variability in assessing EDF levels, we can look at two different time periods as examples. First, looking at the first half of 2009, Figure 12 suggests the EDF value over-predicts. However, looking at Figure 13 makes it clear that default rates peak at levels close to the EDF level peak. As another example, focusing on 2002, on the one hand, Figure 12 suggests that in the first half of the year the model was moving in the opposite direction of default rates, reaching seemingly over-stated EDF levels by the second half. On the other hand, looking at Figure 13, we find multiple months in

2003 with default rates comparable to levels predicted by EDF level. These default rates are masked in Figure 12's filter. In this respect, the model is doing a fine job capturing risk levels in different times.

Second, while the model does not over-predict at the peak of 2009 recession, one can argue that EDF levels are often slightly higher than actual default rates during 2008. It is important to reflect back on the early period of the Great Recession, as financial institutions experienced pressure, the credit markets were squeezed, and the future looked bleak. At first, policy response was unknown, and subsequently, when the quantitative easing policy was announced, its implications were untested. In hindsight, we know the Great Recession, as bad as it was, was not as bad as early predictions indicated.[22] Likewise, a market-based model should be expected to stay true to the *ex-ante* assessments of credit risk and not reflect what one would have predicted if one had perfect foresight.

Third, the data censoring that results from defaulting is not limited to the month of default. It does happen that an exchange delists a firm shortly prior to its default event. To that extent, the default intensity measure can be downward biased.

We remedy this issue by flagging the twelve months leading the default events and looking at the percentage of flagged observations each month. Alternatively, we can flag the twelve months trailing the events. The following table explains the process with an example.

Table 3

Different Definitions of Default Flag

DATE	TIME TO DEFAULT	INSTANTANEOUS	12-MONTH LEADING	12-MONTH TRAILING
2001-03	-15	0	0	0
2001-04	-14	0	0	0
2001-05	-13	0	0	0
2001-06	-12	0	1	0
2001-07	-11	0	1	0
2001-08	-10	0	1	0
...		...		
2002-03	-3	0	1	0
2002-04	-2	0	1	0
2002-05	-1	1	1	0
2002-06	0	-	-	0
2002-07	1	-	-	1
2002-08	2	-	-	1
2002-09	3	-	-	1
...		...		
2003-01	7	-	-	1
2003-02	8	-	-	1
2003-03	9	-	-	1

Table 3 shows a firm that defaulted in June of 2002, as an example of different ways an observed default time series can be constructed. The "12-Month Leading" corresponds with what the EDF measure would be were one to know the date of default 12 or more months ahead of time with perfect foresight. Such a model with perfect foresight has an EDF of 0% until 12 months before default, subsequently hits 100% twelve months before default, and then remains there until default. Although not immune, the 12-month leading definition is less susceptible to data censoring than the instantaneous measure.

The advantage of the 12-month trailing method is in producing results more comparable to the rating agency's statistics. However, since the example in Table 3 was pulled from the stock market nine months after default, constructing a trailing measure of default for a market-based model requires special consideration to account for data-censoring.

Figure 14 shows EDF measures next to the ODF, as defined by the 12-month leading definition. While the model lags the orange line, it does not lag the default event, but rather a hypothetical model with perfect foresight of at least twelve months.

[22] After hitting a bottom in March of 2009, the Dow Jones Industrial Average grew by 60% during the following 12 months and another 15% during the next 12 months.

Figure 14 Time Series of EDF vs. Model with Perfect Foresight

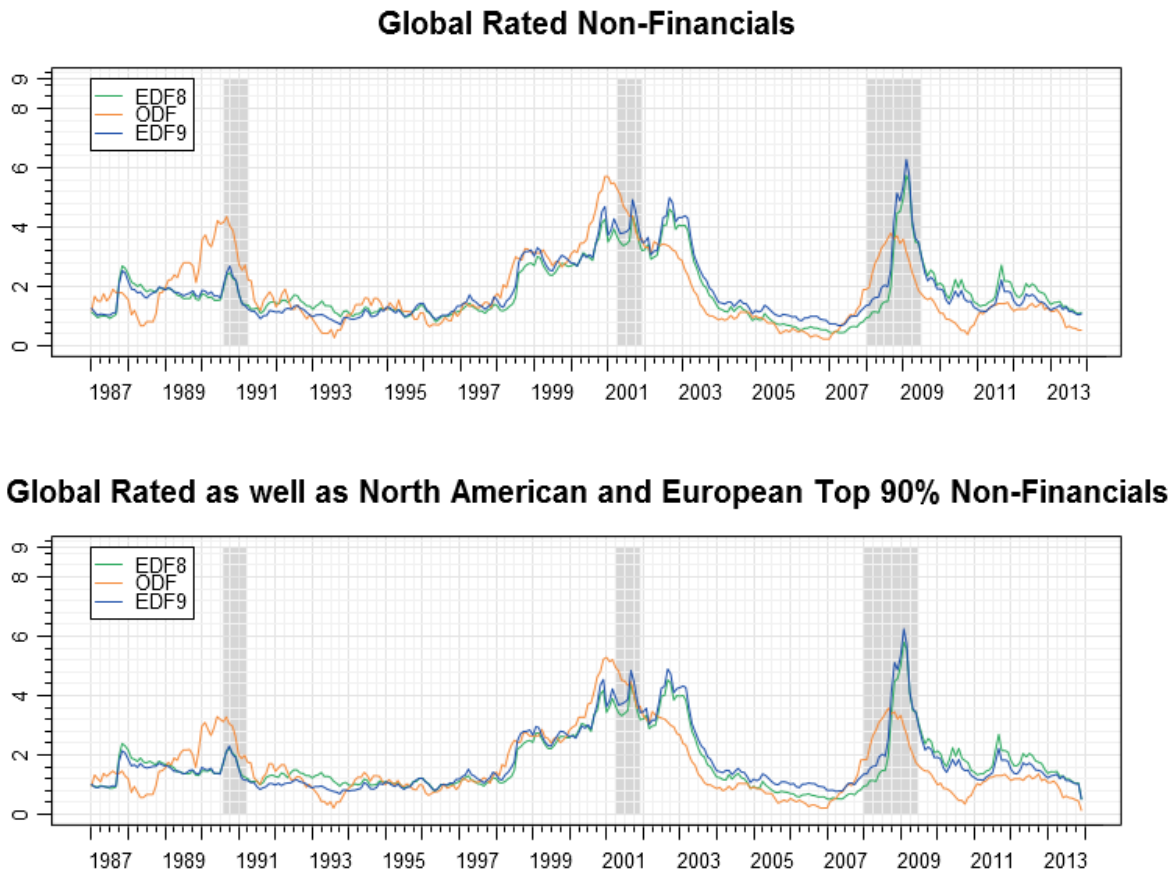

6. Conclusion

Each day, after market close, Moody's Analytics' Public Firm EDF model processes the credit risk of some 35,000 firms globally, producing a default probability named the EDF credit measure. At a very high level, the measure constitutes a Distance-to-Default for each firm and a transformation of the DD into an Expected Default Frequency (EDF).

The DD is built on the causal relationship between default and economic drivers of defaults. It can be heuristically thought of as a risk-adjusted measure of market leverage. The risk adjustment accounts for the business model and its effect on the risk of future cash flows. The market leverage is a ratio of market value of assets, a valuation reflecting net present value of future cash flows, and the default point.

The transformation of DD to physical probability space is such that, at all credit risk levels, the EDF measure is consistent with physical realization of default over an extended period of time. This transformation is calibrated against a global set of firms, spans multiple credit cycles, and is applicable to different macroeconomic environments.

The 9th generation of the Public Firm EDF model is the continuation of the path laid down by KMV in the field of structural credit-risk modeling. The new version continues as a structured option-pricing based model and calculates the put option, i.e. the difference between the risk-free value of liabilities and the risky value of liabilities, using the same formula as its predecessors. The primary model drivers continue to be asset value, default point, and asset volatility.

We have refined the estimation of the drivers and have updated the calibrations. These changes incorporate the experience of the most recent credit cycle, as well as the increase in the number of public firms globally.

The Great Recession deepened our understanding of financial firms. There has been a general understanding that financial firms have a special business model, with special credit risk implications. Given the benefit of more data, we have been able to take this intuition and incorporate it into model implementation. In improving the default point, EDF9 now uses a more parsimonious and transparent formula. This formula has the advantage of increased stability when compared to its predecessor, which, in conjunction with its transparency and its replicability, results in more intuitive movements in firm-level EDF measures for financial firms. Another improvement is with the DD-to-EDF mapping, where a financial-firm specific mapping now results in more prudent EDF levels for financials, as well as more prudent changes to the EDF level in response to observed events. After making all these changes, we see an increase in EDF9's accuracy ratios for financial firms over EDF8.

The measure of asset risk, which adjusts market leverage, has also been improved in several ways. Most notably, benefitting from the global expansion in equity markets, we now model asset volatility at a much more granular level. The asset volatility of each firm is individually calculated using its de-levered equity returns, and that measure is adjusted for sampling variability using our modeled asset volatility. Location is one of the three factors in modeling asset risk, and with EDF9 we now define it at the country level, as neighboring countries within a region may be subject to different asset risk cycles. Furthermore, asset risk is now adjusted to be more relevant to the next year. This results in a more point-in-time measure, also more stable at an aggregate level. Last, we can now compute asset returns using weekly returns over a three-year window globally, whereas, EDF8 used monthly returns over a five-year window outside of North America.

Overall, the most significant improvement is with our ability to address financial firms better. As the EDF metric is now calibrated to the actual failure of financial institutions, stemming from data that now includes the Great Recession period, we have a more accurate measure of the stand-alone risk of such firms. Also, relative to EDF8, the EDF 9 credit metric is more stable for financial firms. The EDF9 metric provides a timely alternative view to the agency ratings of financial institutions.

For non-financial firms, the EDF metric continues to be a valuable early warning risk tool that incorporates the most recent information available on both the firm and the business environment in which it operates.

Appendix A The Public Firm EDF Model within the Context of Moody's Analytics Credit Risk Solutions

Moody's Analytics Public Firm EDF model is the core model behind the product commercially known as CreditEdge and its predecessor Credit Monitor. The core model is also used in a variety of other products and hosts many alternative credit risk models. This section introduces these related products and models.

Each day, the EDF model processes data for a comprehensive set of approximately 35,000 firms, spread globally over almost 90 countries. The resulting statistics, most notable of which is the one-year EDF, are provided to users through multiple channels. One such channel is the website platform creditedge.com. Users can also access model data with an Excel Add-in that allows direct downloading model values into an Excel sheet. Additionally, with the Data File Service (DFS), clients can retrieve data as flat-files via an FTP server. Last, a WebAPI allows for programmatic access.

Beside the one-year EDF, CreditEdge delivers alternative measures of credit risk, using models that build upon the basic Public Firm EDF model. These include the Term Structure model, the Valuation model, the Stressed EDF model, and the Through-the-Cycle model. We also provide a product feature that enables estimating a firm's rating based upon its EDF credit measure — an EDF-implied rating. Finally, our RiskCalc product uses *Distance-to-Default* to adjust the default risk of private firms for the current stage of the credit cycle.

The Term Structure model estimates a term structure of EDF values to measure credit risk at different horizons, between one to 10 years. While, in the short-term, a firm is subject to both idiosyncratic and systematic risks, over longer periods of time, spanning the duration of a business cycle, the idiosyncratic risks dominate, giving the credit risk measure acyclical properties. Such long-term EDF measures can be used for valuing long-duration exposures. Longer-term EDF metrics are more stable, which may make them more suitable for capital requirement calculations.

The EDF-implied rating estimates a rating for every non-rated firm in the CreditEdge database, as implied by its EDF metric, as well as the nine EDF values on its term structure. The relationship between EDF credit measures and agency ratings can be calibrated based on either Moody's ratings or Standards & Poor's ratings. This feature enables an estimate of what a firm's ratings would have been, had it been rated. For rated firms, the EDF-implied rating enables studying the difference between the market sentiment, as reflected through a structural quantitative model and that of fundamental analysis done by rating agency analysts. Furthermore, as the EDF metric often responds faster to changing market conditions, the EDF-implied rating may also be used to speculate on rating upgrades and downgrades.

While the EDF measure reflects information embedded in the equity markets, the Valuation model can be used to better understand the credit risk priced into credit default swaps (CDS) or bonds. The Valuation model also enables estimation of the CDS implied by the EDF measure. This information can be useful in both mark-to-market valuation of illiquid instruments, as well as in developing trading strategies.

The Stressed EDF measures model credit risk conditional on macro-economic factors. In other words, we can devise a hypothetical macroeconomic scenario and then study the effects on firm-level credit risk by producing conditional EDF measures. Stress testing is the most common use case of Stressed EDF measures. The Through-the-Cycle EDF model produces EDF measures that have been filtered to remove the credit cycle. The TTC EDF metrics reflect the firms' underlying long-run credit risk.

Another related product is Moody's Analytics RiskCalc, RiskCalc produces private firm EDF measures to produce a forward-looking default probability by combining financial statement and a credit cycle adjustment into a highly predictive measurement of stand-alone credit risk for private firms. The credit cycle adjustment uses the Moody's Analytics Public EDF model to adjust the risk of private firms to reflect the recent trends in public firms in comparable regions and industries.

All the related and extended EDF models produce the credit risk of a single obligor. When many exposures are aggregated into a portfolio, quantifying the resulting portfolio risk poses another set of challenges. Tackling that problem requires knowledge of exposure correlations. For that purpose, Moody's Analytics offers the Global Correlation Model[23] or GCorr. The GCorr model uses asset returns derived from the core Public EDF model to estimate correlation of different firms. RiskFrontier calculates the loss distribution of a portfolio, where key inputs are EDF metrics and correlations.

[23] For more details please refer to Huang (2012).

Appendix B Conceptual Grounding of the Public Firm Model

This section is intended for those new to the Public Firm model framework. It conveys the conceptual grounding of the model by taking a balance sheet approach, beginning with leverage.

We seek to build models with a solid intuitive foundation that perform well and are applicable. Critical to model application is the assessment of the reasons behind output change and without a transparent and intuitive structure such assessment is out of reach. Since the thought process behind the Public Firm EDF model is one of studying the causal relation between default and the economic drivers of default, there is a very grounded intuition behind its structure. At the heart of the model is a rank-ordering statistic, Distance-to-Default (DD), which can be thought of as an evolution of a balance sheet-based measure of leverage. This section explores the intuition behind DD. To explain this intuition, we incrementally refine the balance sheet-based measure of leverage, in four steps. In each step, we describe a refinement and illustrate its value using an example. We then arrive at a heuristic definition of DD.

The balance sheet measure of leverage we begin with is the ratio of book equity to book value of total assets, perhaps one of the oldest tools in credit analysis. At its core, the Public Firm EDF model is not a different paradigm, but rather a modern approach to measuring this value. Indeed, the result is a credit metric we can describe as risk-adjusted market leverage. In other words, the measure can be thought of as a leverage, measured with market valuations, and adjusted for asset risk. Through the next set of examples, we consider each in turn.[24]

Table 4 summarizes the four definitions of leverage we consider, each an incremental refinement of the prior.

Table 4

Different Leverage Definitions

USING	
Book values	$\frac{Book\ Equity}{Book\ Value\ of\ Total\ Assets}$
Equity market values	$\frac{Market\ Capitalization}{Market\ Capitalization + Total\ Liabilities}$
Market Value of Assets	$\frac{Market\ Value\ of\ Assets - Liabilites}{Market\ Value\ of\ Assets}$
Term structure of liabilities	$\frac{Market\ Value\ of\ Assets - Default\ Point}{Market\ Value\ of\ Assets}$

For the first example, we compare two firms, Eastman Kodak and Cablevision System Corporation. We find that six months before default, Eastman Kodak had a leverage ratio of -22%, indicating that book liabilities exceeded book assets. Kodak was an early pioneer in film photography, but Kodak was unable to successfully adapt its business model to the innovation of digital photography. Nevertheless, companies can have a negative net worth for a variety of reasons and yet still maintain an economically viable business. Indeed, in the same month, Cablevision Systems Corp. had a leverage ratio of -72%, yet it has not defaulted to date. The relevant reason was not reflected in the balance sheet, but it was accounted for by the equity market. The market capitalization of a company reflects the markets' perception of the future earnings of a company. Consequently, market capitalization is more forward-looking than book equity, and, in our view, it is more useful for understanding the credit risk of a company. If we define leverage as market capitalization as a percentage of the firm value, where firm value is market cap plus total liabilities, we find Kodak's leverage ratio was 8% and Cablevision's 31%. In other words, by using market capitalization, we can actually rank-order Kodak as riskier than Cablevision.

[24] We study each pair six months before the default of the defaulting member of the pair. In these examples, the financial statement is actually one quarter prior to the market date to allow for time for both the firm to finalize the statements, as well as time for the vendors to process the statement into a data feed. In one case, there is a two quarter lag, as that firm only produces semi-annual statements.

Table 5

Book Value of Equity vs. Market Capitalization

	EASTMAN KODAK	CABLEVISION SYSTEM CORPORATION
Book Equity	-1,274	-6,462
Book Assets	5,882	8,963
Market Capitalization	646	7,008
Total Liabilities	7,156	15,425
Book Equity / Book Assets	-22%	-72%
Market Cap / (Market Cap + Total Liabilities)	8%	31%
unit	USD MM	USD MM
statements	2011-03	2011-03
market price: last trade of	2011-07	2011-07

We have now made the first refinement to our leverage measure. For the next refinement, we improve our value of assets estimation. Firm value is one estimate of the value of a firm's assets. A second estimation is that made by an accountant, where total assets is the summation of what is paid for when purchasing a company's assets, after accounting for depreciation and other accounting rules. So at the beginning, the total value of a firm's assets is the total funding provided by the right hand side of the balance sheet. A third estimation comes from a market valuation perspective. In this case, the value of assets comes from the net present value of the future cash flows they generate, reflective of the synergies of the different assets, not a simple summation of them. Therefore, the market value of assets can differ from the value of assets reported in the statements or from the firm value.

The market value of assets has two components:

$$Market\ Value\ of\ Assets = Market\ Value\ of\ Equity + Market\ Value\ of\ Liabilities$$

The market value of a firm's liabilities differs from the book value of liabilities. The book value of liabilities is typically the discounted value of promised future payments, discounted at the loan's coupon. As a firm enters financial distress, the market value of liabilities typically declines, as one could buy the liabilities for less than the book value of the debt. This decline reflects the increasing default risk premium that investors demand in order to hold the risky debt. Therefore, the market value of assets for a firm in financial distress is typically less than the market capitalization plus the book value of total liabilities. The Public Firm EDF model estimates this default risk premium and uses it to obtain the market value of assets, a concept similar to Enterprise Value.[25]

We can see the benefit of using the market value of assets in the next example. Six months before its bankruptcy, Lehman Brothers Holdings leverage ratio was 3%, using firm value. However, accounting for the default risk premium of Lehman's liabilities, we find its leverage ratio was -15%.

[25] The key difference between enterprise value and what we call the market value of assets is the treatment of cash. In the context of acquisitions, the enterprise value of a firm is intended to estimate how much it costs to buy the company debt-free and cash-free. In the context of the our model, the market value of assets is how much it costs to buy the company debt-free.

Table 6

Market Cap + Total Liabilities vs. Market Value of Assets

	LEHMAN BROTHERS HOLDINGS	BARCLAYS PLC
Market Capitalization	19,971	29,752
Total Liabilities	637,483	1,129,283
Default Risk Premium	105,533	100,611
Market Value of Assets	551,921	1,058,424
Market Cap / (Market Cap + Total Liabilities)	3%	2.5%
(MV of Assets – Total Liabilities) / MV of Assets26	-15%	-7%
unit	USD MM	GBP MM
statement	2007-08	2007-06
market price: last trade of	2008-03	2008-03

At the same time, we see that Barclays' credit risk is measured as less risky than Lehman's, only after proper estimation of the default risk premium and the subsequent estimation of market value of assets.

For the third refinement to leverage, we look at total liabilities. Is total liabilities indeed the best measure to compare against asset value? Relative to current liabilities, long-term liabilities provides firms some breathing room to work through transitory disruptions in cash flow. A firm with more immediate due dates is at higher risk than one with due dates years into the future. Consequently, the two measures are not of the same importance in determining default risk at the one-year horizon.

Take Bombardier for example. Six months before default, its leverage ratio was 22%, as was Bouygues SA. Nevertheless, Bouygues did not default (as of yet). Backed by years of research, we find that the default point is a more suitable value to compare against asset values. Indeed, when we define leverage using the default point, we find Bombardier's leverage to be less than Bouygues'. The so-called default point is defined as current liabilities plus one-half long-term liabilities. Put differently, this study tells us that, while the market value of assets was dropping, Bombardier had a relatively large portion of its liabilities due in less than a year, and this particular capital structure may have contributed to its default.

Table 7

Total Liabilities vs. Default Point in Construction of Leverage

	BOMBARDIER	BOUYGUES SA
Market Capitalization	5,403	9,174
Default Risk Premium	40	1,244
Short-Term Liabilities	14,880	18,836
Long-Term Liabilities	3,873	8,907
Default Point	16,816	23,289
Market Value of Assets	24,116	35,673
(MV of Assets – Total Liabilities) / MV of Assets	22%	22%
(MV of Assets – Default Point) / MV of Assets	30%	35%
unit	CAD MM	Euro MM
statement	2008-10	2008-09
market price: last trade of	2009-01	2009-01

We began with a basic measure of leverage and in three steps refined it to use market data, incorporate a default risk premium, and take into consideration the relative different importance of short-term and long-term liabilities.

To motivate the final refinement, we now ask if this analysis is even fair. Can we legitimately compare the leverage ratio of a Canadian airplane and train manufacturer to the leverage ratio of a French construction company? Producing different products, they have different cash flow risks. Involvement with different industries subjects them to different regulatory risks, and, additionally, their assets are subject to different monetary and fiscal policy shocks. These considerations, and many more, affect the firm's cash flow risk.

[26] It is noteworthy that the leverage measure is not defined as market cap as a percentage of market value of assets. The reason: the chosen ratio penalizes high credit risk names more than a simple market cap to assets ratio would:

$$Market\ Value\ of Assets = Market\ Cap + Total\ Liabilities - Default\ Risk\ Premium$$

$$\frac{Marekt\ Value\ of\ Assets - Total\ Liabilities}{Market\ Value\ of\ Assets} = \frac{Market\ Cap - Default\ Risk\ Premium}{Market\ Value\ of\ Assets}$$

A proper rank-ordering of credit risk cannot be achieved without accounting for asset risk. If a bank and a pharmaceutical company have the same market leverage, it does not imply they have the same credit risk, as a bank's assets are much safer than a pharmaceutical company's. If we measure the fluctuations of each firm's asset returns we can estimate asset volatility as the standard-deviation of asset returns.

If we take our refined measure of leverage and adjust it for asset volatility, we get

$$\frac{Market\ Value\ of\ Assets - Default\ Point}{Market\ Value\ of\ Assets\ \cdot Asset\ Volatility}$$

Japan Airlines filed for bankruptcy in January of 2010. Six months earlier, the difference between its market value of assets and default point, as a percentage of value of assets, was 44%, while the same ratio for Nagoya Railroad, a ground transportation business in Japan of comparable size was 42%. The default risk premium in both cases is comparable to their market capitalization. In fact, from this angle, Nagoya's situation appears a bit worse. When looking at their asset risk, however, we find Japan Airline's asset volatility was 9%, higher than Nagoya's 6%. A business model working with high-risk assets, for the chosen capital structure, ultimately jeopardized the business.

Table 8

Market Leverage vs. Risk-Adjusted Market Leverage

	JAPAN AIRLINES	NAGOYA RAILROAD
Market Capitalization	440	277
Default Risk Premium	413	282
Short-Term Liabilities	679	483
Long-Term Liabilities	970	463
Default Point	1,164	714
Market Value of Assets	2,062	1,228
Asset Volatility	9%	6%
(MV of Assets – Default Point) / MV of Assets	44%	42%
(MV of Assets – Default Point) / MV of Assets * Asset Volatility	4.8	6.9
Unit	Billions of JPY	Billions of JPY
Statement	2008-12	2008-12
market price: last trade of	2009-07	2009-07

As our risk-adjusted measure of market leverage indicates, Japan Airlines was indeed riskier than Nagoya during the months preceding the credit event.

We have now constructed a heuristic version of DD using market value of assets, default point, and asset volatility. This ratio effectively describes the distance of the firm from default in standard deviation units. In the context of our model, the default point is the point where creditors lose faith in the firm's ability to service its current obligation, stop extending credit, and push the firm into default. So the point where the market value of assets reaches the default point is important. On the other hand, as time goes by, what regulates asset values' ups and downs is asset volatility. Given leverage, the more volatile the assets, the more likely it is that their value crosses the default point in the future. In other words, what we are interested in is the gap between the market value of assets and the default point, as a percentage of asset value fluctuations, and that is precisely what our updated leverage measure shows.

It is at first surprising that a single model can produce such reliable measures of credit risk for more than 35,000 different firms, across such varied countries and business activities. But part of the model's effectiveness in rank-ordering is in separately calculating the market leverage of each firm, as well as its asset risks; coming up with an individualized DD that properly controls for the factors relevant to credit risk.

While the tables above reported ratios in percentages, the last table reports the DD of JAL and Nagoya as simply 4.8 and 6.9, respectively. The units are standard deviations, so they are not yet useful for computing the expected default frequency of an individual name or the expected loss rate of a portfolio. To rectify this issue, we transform the DD to the physical probability space, and call it the EDF credit metric. This process is further described in Section 4.

The Public Firm EDF model measures a specific aspect of credit risk — the physical probability of default — contrasted with credit default swaps, whose spread reflects loss given default, as well as the risk-neutral probability of default. Moreover, EDF measures

reflect a firm's stand-alone credit risk without the aid of government support (e.g. a government bailout). The EDF credit measure is a scaled metric, meaning its absolute value has an objective interpretation. Consequently, the EDF credit measure has cyclical properties, thereby reflecting changes in systematic default risk. On the other hand, ratings published by Moody's Investors Service provide "relative credit worthiness" (page 2, Moody's Investors Service, 2004) as an ordinal measure. An agency rating may remain stable through a credit cycle, yet the default risk of a rating, for instance a Baa3, will be cyclic. Furthermore, as the market value of equity is a primary driver, EDF value is not directly impacted by the liquidity of the bond market or the CDS market.

Appendix C Calibration Sample Size Cuts

This appendix explains the size cuts used to construct different samples against which we test the DD to EDF calibration.

In order to capture the population of top-90%, once a month and in each region, we sort firms by size from large to small, and choose until we have 90% of the outstanding total liabilities of that region for that month. In other words, the model is calibrated to where 90% of credit risk resides, as well as for rated firms.

Figure 15 Size of the 90%, 99%, and 99.9% Cuts, Non-Financial Corporations

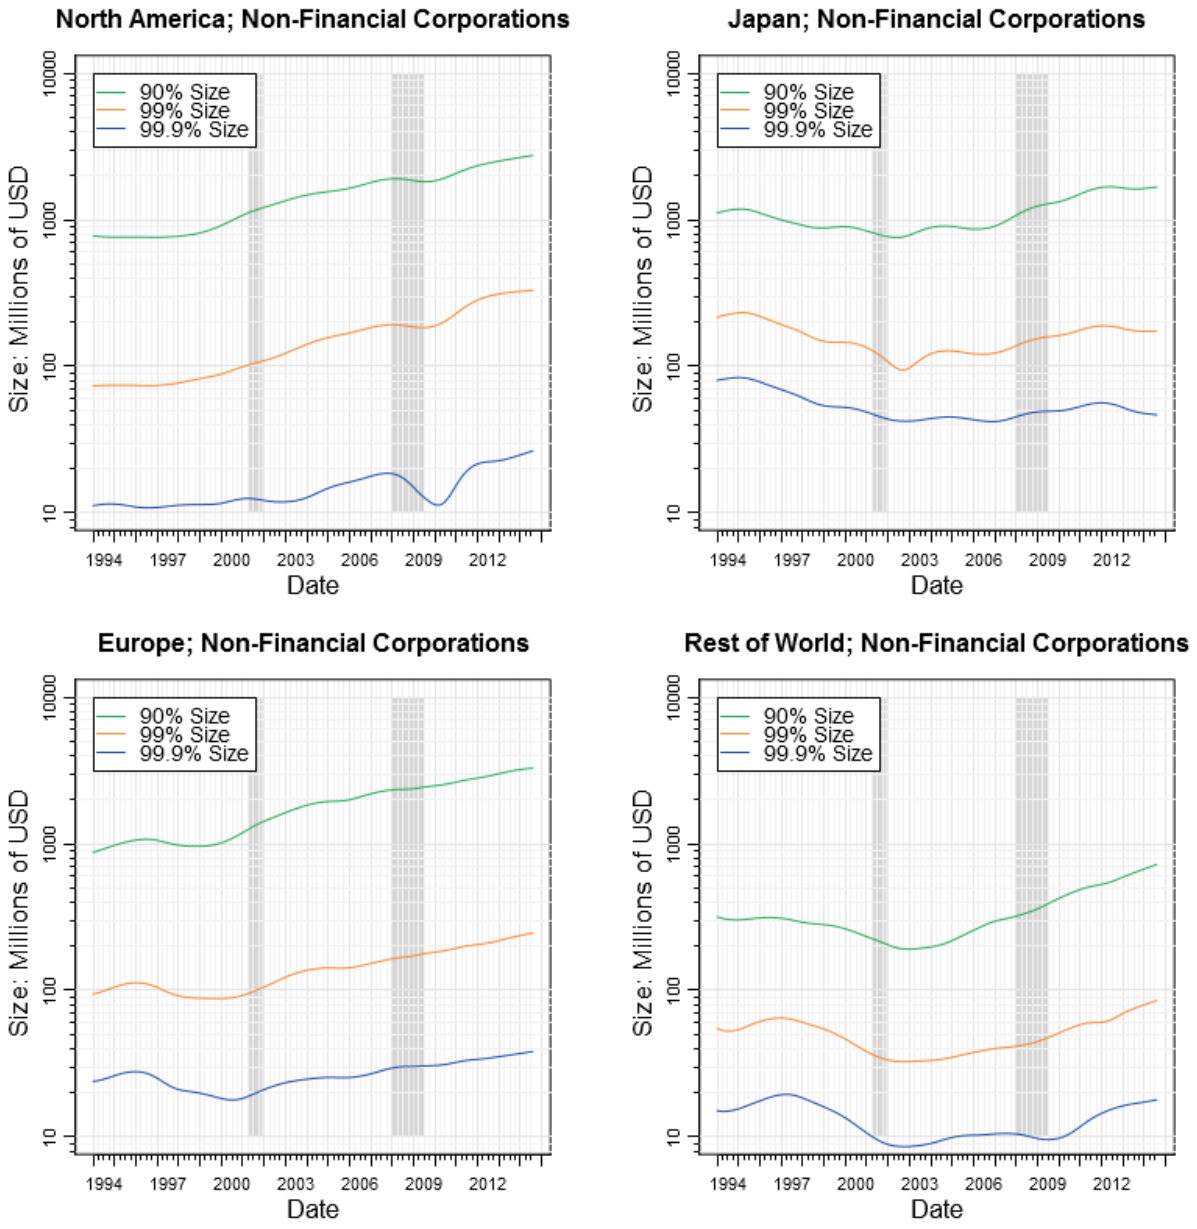

Figure 16 Size of the 90%, 99%, and 99.9% Cuts, Financial Firms

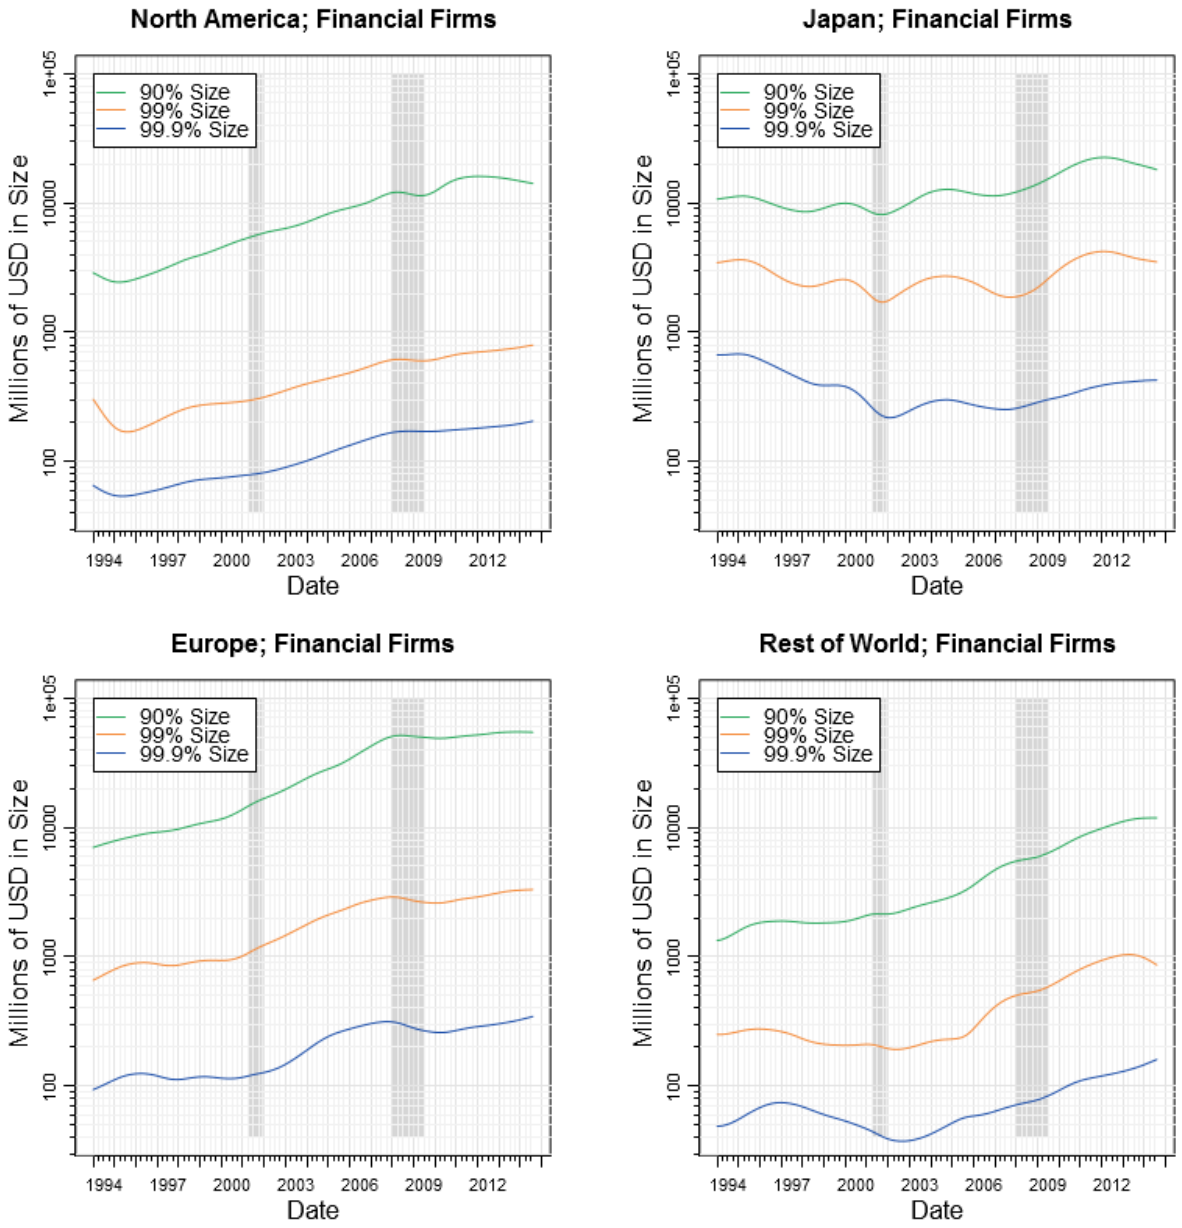

For non-financial corporations, the top-90% cut-off roughly coincides with a size of about two billion dollars. For financials, the cut-off is closer to twenty billion dollars. The two figures above show the relation between size and the 90%, 99%, and 99.9% cut-offs, where size is the average of book value of assets and sales.

References

Arora, Navneet, Jeffery Bohn, and Irina Korablev, "Power and Level Validation of the EDF Credit Measure in the U.S. Market." Moody's KMV, 2005.

Arora, Navneet, Jeffery Bohn, and Fanlin Zhu, "Reduced Form vs. Structural Models of Credit Risk: A Case Study of Three Models." *Journal of Investment Management,* Vol.3, No.4, Fourth Quarter 2005.

Basel Committee on Banking Supervision, "International Convergence of Capital Measurement and Capital Standards ('A Revised Framework')." Bank for International Settlements, 2004.

Black, Fischer and John C. Cox, "Valuing Corporate Securities: Some Effects of Bond Indenture Provisions." *The Journal of Finance* 31:351-367, 1976.

Bohn, Jeffrey R. and Roger M. Stein, *Active Credit Portfolio Management in Practice*. Vol. 384. John Wiley & Sons, 2009.

Cantor, Richard and Chris Mann, "Analyzing the Tradeoff Between Ratings Accuracy and Stability." Moody's Investors Service, September 2006.

Caouette, John B., Edward I. Altman, Paul Narayanan, and Robert Nimmo, *Managing Credit Risk: The Great Challenge for Global Financial Markets*. Vol. 401. John Wiley & Sons, 2011.

Chen, Nan, Houman Dehghan, Min Ding, Jian Du, James Edwards, Danielle Ferry, Pooya Nazeran, Sue Zhang, Douglas Dwyer, and Jing Zhang, "EDF9: Introduction Overview." Moody's Analytics White Paper, February 2015.

Crosbie, Peter and Jeff Bohn, "Modeling Default Risk." Moody's KMV White Paper, December 2003.

Duffie, Darrell and Kenneth J. Singleton, *Credit Risk: Pricing, Measurement, and Management: Pricing, Measurement, and Management*. Princeton University Press, 2012.

Dwyer, Douglas and Shisheng Qu, "EDF 8.0 Model Enhancements." Moody's KMV, January 2007.

Engelmann, Bernd, Evelyn, Hayden, and Dirk Tasche, "Testing Rating Accuracy." *Risk*, January 2003.

Huang, Jummy, Mariano Lanfranconi, Nihil Patel, and Libor Pospisil, "Modeling Credit Correlations: An Overview of the Moody's Analytics GCorr Model." Moody's Analytics White Paper, December 2012.

Kealhofer, Stephen, "Quantifying Credit Risk I: Default Prediction." *Financial Analysts Journal,* January/February 2003, 30-44.

Kealhofer, Stephen and Matthew Kurbatt, "Predictive Merton Models." *Risk*, February 2002.

Knowledge Base FAQ 388, "Estimation of Short- and Long-Term Liabilities." Moody's KMV, May 2006

Knowledge Base FAQ 438, "What Goes into Asset Volatility and Asset Value Calculations?" Moody's KMV, April 2009.

Korablev, Irina, "Power and Level Validation of the EDF Credit Measure in the European Market." Moody's KMV 2005.

Leland, H.E, "Corporate Debt Value, Bond Covenants, and Optimal Capital Structure." *The Journal of Finance* 49(4), 1994.

Miller, Ross, "Refining Ratings." *Risk*, August 1998.

Moody's Analytics, "Public EDF Validation." Moody's Analytics, March 2015.

Ranson, Brian J., *Credit Risk Management*. Sheshunoff/Alex eSolutions, 2005.

Stein, Roger, "Evidence on the Incompleteness of Merton-Type Structural Models for Default Prediction." Moody's KMV, 2005. This paper is a revised version of the Moody's Risk Management Services paper from 2000.

Sun, Zhao, David Munves, and David Hamilton, "Public Firm Expected Default Frequency (EDF) Credit Measures: Methodology, Performance, and Model Extensions." Moody's Analytics White Paper, 2012.

Vasicek, Oldrick, "A Note on Using Cross-Sectional Information in Bayesian Estimation of Security Betas." *Journal of Finance,* Vol. 29, May, 449-470, 1974.