

Addendum to *Public Firm Expected Default Frequency (EDF™) Measures – At a Glance*

Overview

This addendum to “Public Firm Expected Default Frequency (EDF™) Measures – At a Glance” provides additional detail on the derivation of the public firm EDF model drivers and the EDF-implied ratings methodology. These descriptions should be sufficient for most readers seeking a working understanding of the public firm EDF model and EDF-implied ratings. Those seeking in-depth model methodology, to satisfy internal- or regulatory-driven model validation requirements, should refer to Nazeran and Dwyer, “Credit Risk Modeling of Public Firms: EDF9,” “Moody’s Analytics Modeling Methodology,” June 2015.

EDF Model Drivers

Asset Value

A company’s market value of assets is not directly observable, but its equity value is. Asset value is the present value of all future cash flows generated by the assets, while equity value is the present value of the portion of total cash flows received by equity holders (after the company’s debts are serviced). This indirect link between a company’s equity value and its asset value allows us to estimate the latter.

One of the key insights underpinning the EDF model is the observation that equity holders possess a walk-away option. If the value of a firm is lower than the value of its liabilities, shareholders can and will let creditors assume ownership of the firm. Owning the equity of a company, therefore, is analogous to holding a call option on the company’s assets. A call option gives an investor the right to buy a specified stock (or other asset) at a specified (strike) price within a specified time period. The investor benefits when the price of that stock rises above the strike price, but loss is limited to the premium of the option. Similarly, an equity holder’s loss is limited to the acquisition cost of the shares. In our equity holder analogy, the required debt payment at the horizon date serves as the option’s strike price.

The classic Black-Scholes model provides a formula for calculating the unknown value of an option using the observable value of the underlying asset (stock price) and other inputs. In the case of the EDF model, the value of the option is observable (equity value), and we seek to uncover the value of the underlying asset (asset value) from this and asset volatility. Thus, we solve backwards from the option price (market capitalization) to derive each firm’s implied asset value.

In practice, the EDF model diverges from the theoretical Black-Scholes model to introduce more realistic assumptions. This provides a better approximation of real-world firm capital structures and firm behavior. As an example, while the basic structural model assumes default occurs only at the horizon date (i.e., the maturity of the liabilities), the EDF model allows for default at any time.

Default Point

The default point represents the amount of a firm's liabilities due at a given horizon that would trigger a default if not paid according to the contractual terms. In the basic structural model described earlier, the determination of the default point is straightforward – it is simply the face value of a zero-coupon bond. In the practical world of the EDF model, firms' liabilities are often comprised of multiple classes of debt with various maturities, and potential default events are not limited to a specific maturity date. This means that the same amount of total liabilities for two otherwise identical firms may have different required debt payments for a given time horizon. So, the default point is a function of both the prediction horizon and the maturity structure of the liabilities.

The EDF model strives to find a parsimonious representation of firms' default points for a given time horizon (i.e., one that uses a small number of inputs), while maximizing the model's default prediction power. Analysis of our large corporate default database shows that firms do not always default when their asset values fall to the level of their total liabilities – some are able to stay afloat even when they have difficulty meeting their long-term obligations. On the other hand, some firms default before their assets values fall to the book value of their short-term liabilities. In distress, many of a firm's long-term liabilities become short-term liabilities as covenants are broken and debt is restructured.

Partly to address these factors, we employ separate calculations for non-financial and financial firms. For non-financial firms, the default point for a one-year horizon is set at 100% of short-term liabilities plus 50% of long-term liabilities. For financial firms, the default point is set at 75% of total liabilities. While deceptively simple, these rules produce exceptionally accurate EDF measures, particularly when we add in the cost-of-capital adjustment described below.

As a last step in the default point calculation, we adjust for cost-of-capital. Multiplying the default point by $1+r$ (the risk-free interest rate) improves upon a model that would otherwise tend to underestimate credit risk during high interest rate environments and overestimate credit risk in low interest rate environments. The intuitive explanation for this is that when interest rates are high, companies typically find it more difficult to roll over credit and to accumulate unpaid bills. Working capital depletes faster, and access to credit becomes more expensive.

Asset Volatility

Like the market value of assets, asset volatility is not observable and must be estimated. Note that while the two are related, asset volatility is distinct from equity volatility. Moreover, the asset volatility of a firm is smaller than its equity volatility. This observation makes economic sense, since equity holders receive cash flows only after the company's debt is serviced. Therefore, their cash flows exhibit greater volatility than for the company as a whole (i.e., for its assets).

To produce stable and informative estimates of a company's asset volatility, we incorporate both firm-specific asset return information and information from comparable firms. The volatility estimate based exclusively on a firm's own asset return information is termed *empirical volatility*, and the volatility estimate derived from the information of comparable firms is termed *modeled volatility*. The combination of the two, plus a *forward-looking volatility adjustment*, yields the total asset volatility measure used in the EDF model.

For firms with sufficient data, we calculate *empirical volatility* using a three-year window of weekly asset returns. Since asset returns are also not observable, we must estimate them. Here, we rely on the equality of the market value of assets to the market value of equities plus the market value of liabilities. In this way, asset returns can be calculated as the sum of de-levered equity returns and the returns on long-term debt multiplied by the ratio of long-term debt to asset value. A key input, then, to estimated asset returns and thus asset volatility, are weekly observations of companies' stock prices. Before calculating empirical volatility, we make adjustments to asset returns to account for large corporate events such as mergers or acquisitions, and we trim outliers.

Modeled volatility can be thought of as the empirical volatility of an average firm with a certain set of characteristics. These characteristics are defined by firm size, industry, and geography. We use 61 industry groups and 68 country groups. Since the systematic asset risk embedded in each firm's country-industry group is cyclical, we recalibrate modeled volatility each month so that on average, modeled volatility is equal to empirical volatility.

Combined asset volatility is obtained by taking a weighted average of empirical and modeled volatility, where the weight on empirical volatility is an increasing function of the availability of asset returns. For new firms, only modeled volatility is used. Over time, as more weeks of asset returns become available, the model gradually places more weight on empirical volatility. The weight on empirical volatility bottoms at 20%, even for firms with more than 3 years of weekly asset returns.

There are a number of benefits associated with incorporating modeled volatility into the asset volatility calculation. First, it enhances the stability of asset volatility estimates by reducing the noise contributed by firm-specific returns. Second, it improves the forward-looking default prediction power of the resulting EDF measures.

A final adjustment to the combination of empirical and modeled volatility further enhances the forward-looking power of the EDF model. The *forward-looking adjustment* accounts for mean reversion observed in our data: A three year period of relatively high volatility is followed by lower asset volatility in the following year, and vice versa. Therefore, we adjust volatility upward when it is low on a historical basis, and we adjust volatility down when it is high on a historical basis. This results in a more accurate, yet stable, measure of credit risk.

Distance-to-Default (DD)

The three key drivers in the EDF model – asset value, the default point, and asset volatility – combine to form distance-to-default (DD). DD, which is approximately equal to market leverage divided by asset volatility, represents the distance between the expected market value of assets at the horizon date (A , in the formula below) and the default point (X), standardized by its business risk in the form of its asset volatility (σ_A).

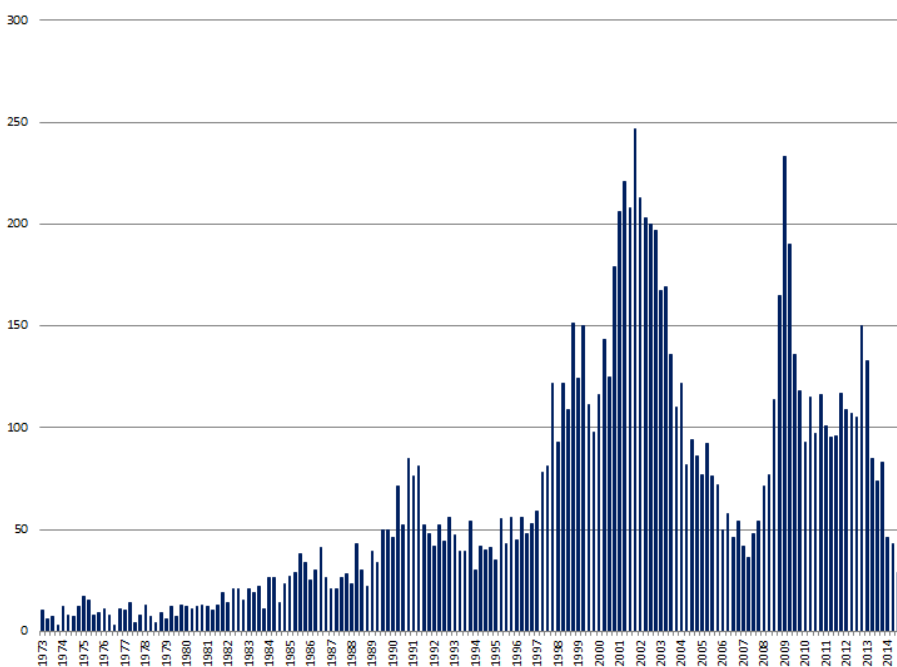
$$DD \approx \frac{\ln\left(\frac{A}{X}\right)}{\sigma_A}$$

Distance-to-default provides an effective rank ordering statistic of default risk. We observe a strong, negative empirical relationship between DDs and observed default rates. This makes sense: a firm is more likely to default if the gap between its assets and liabilities is small (i.e., its financial leverage is high) and its asset volatility (business risk) is elevated. However, DDs are ordinal measures, that is, they provide a rank order estimate of default risk, but not the precise level of risk. Probabilities of default (PDs), by their nature, provide exact risk rates, and thus are cardinal measures. Therefore, one additional step is needed in order to generate a PD.

Moving from DD to EDF

Basic structural credit risk models tend to overestimate default risk at very small DD levels and to underestimate default risk at very large DD levels. To avoid this issue, we employ an empirical mapping of DD to EDF (i.e., one based on historical default data), which leverages Moody's Analytics comprehensive default database of over 11,700 global defaults between 1973 and 2014. Figure 1 plots these defaults over time. The breadth of our default database, both in terms of geographical reach and time (spanning five US economic cycles) permits a finely tuned calibration of DD to EDF and is one of the model's key advantages over other PD models.

Figure 1 Historical Default Database (Count of Defaults Per Year), 1973-2014



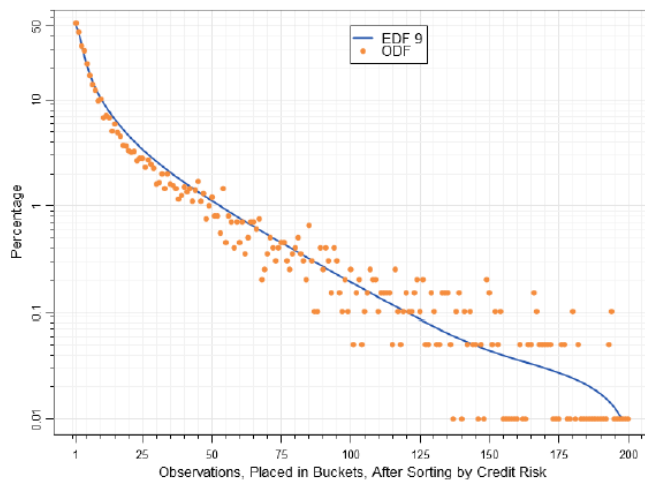
The calibration relies on the long-term relationship between DDs and observed default rates. Since these relationships are quite different for financial firms than for non-financial firms, we employ two separate DD-to-EDF mappings. Figure 2 illustrates how the calibrations are constructed and shows why separate calibrations are necessary. Firms are first grouped into equally-sized buckets according to their DD levels, from the lowest on the left to the highest on the right. These DD level buckets are represented on the charts along the horizontal axis. Next, using our default database, we calculate observed default frequencies (ODFs) for each bucket (i.e., the percent of firms in each bucket that defaulted over rolling one-year time horizons). These are represented on the charts by the orange circles. Finally, we fit a nonlinear function to the relationship between DDs and ODFs.

As Figure 2 makes evident, the relationship between DDs and ODFs is quite different for financial firms than for non-financial firms. In particular, the relationship is much flatter for financial than for non-financial firms – a given change in DD bucket translates into a smaller movement in actual default risk for financial firms. This is intuitively consistent with the less transparent nature of financial firms' financial statements compared with those of non-financial companies, which reduces the signal-to-noise ratio of the market signals. The separate calibrations contribute to more stable EDF measures through time, particularly among financial firms.

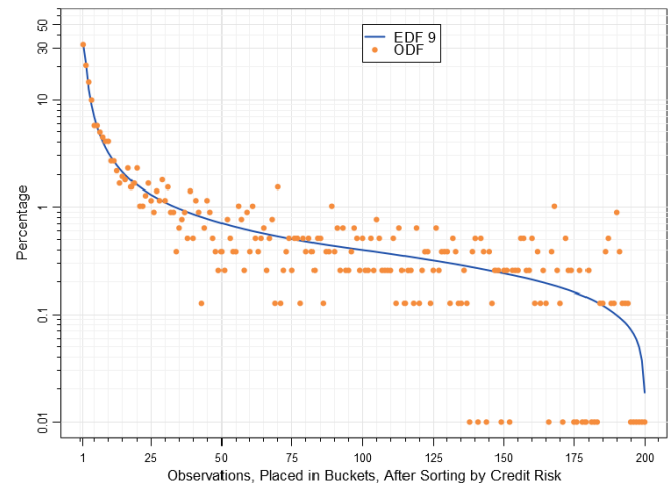
The charts in Figure 2 also shed light on the caps we impose on EDF measures. While EDF measures for non-financial firms are capped at 50%, they top out at 35% for financial firms. These caps are dictated by the data. That is, at very low DD levels, there are enough defaults to differentiate risk up to 50% probability of default (or 35% for financials), but we cannot reliably say, for example, that one firm's probability of default is 51% whereas another's is 60%.

Figure 2 Average EDF Level vs. Observed Default Frequency

(a) Global Rated Non-Financial Firms



(b) Global Rated Financial Firms



EDF-Implied Ratings

CreditEdge users sometimes prefer rating scales as measures of credit risk, rather than probabilities of default. To a large degree this reflects the widespread use of such scales, but the practice is also common because many regulators require a mapping of internal ratings to ratings grades with sufficient risk differentiation. Therefore, we map EDF metrics to the Moody's and S&P rating scales as well as to users' own internal rating scales. CreditEdge has implied ratings for over 42,000 firms, including those without ratings. We utilize separate mapping tables for non-financial and financial firms.

Figure 3 Sample EDF-Implied Rating Table

Implied Rating Table **Internal Rating Table**

The following table displays the correlation between EDF9 Implied Ratings and Median EDF9 credit measures. These values are updated on a monthly basis.

Source: EDF9 Median Calculation: EDF9 Implied Rating: IR Mapping Date:

EDF9 Implied Rating	Entity Type	Equivalent Median EDF9 Credit Measure	Bound Range	
			Greater Than or Equal To	Less Than
Aaa	Corporates	0.0100%	0.0100%	0.0126%
Aaa	Financials	0.0446%	0.0100%	0.0562%
Aa1	Corporates	0.0159%	0.0126%	0.0178%
Aa1	Financials	0.0707%	0.0562%	0.0794%
Aa2	Corporates	0.0200%	0.0178%	0.0214%
Aa2	Financials	0.0891%	0.0794%	0.1001%

EDF-Implied Ratings provide the following benefits:

- » They have a consistent meaning across industries and geographies.
- » Over the long run, the distribution of EDF-Implied Ratings is consistent with distribution of ratings assigned by Moody's Investors Service.
- » They offer the ability to assign implied ratings to unrated firms.

EDF-implied rating mappings are constructed from two samples of firms – rated non-financial firms and rated financial firms. Then, for each sample, we calculate the observed median EDF for each “major” rating category (e.g., Aaa, Aa, A, etc.). We interpolate median EDF values for “minor” rating categories (e.g., Aa1, Aa2, Aa3, etc.) by assigning the median EDF value of the major category to the middle of the corresponding minor category (e.g., Aa and Aa2) and taking a weighted geometric mean of the medians from each rating category's major neighbors. For example, to derive the median EDF value for Aa3, we take the weighted geometric mean of its neighbors, Aa2 and A2.

We impose a minimum median EDF distance between major rating category neighbors, however, so as to limit implied ratings compression, which can occur during benign credit periods when EDFs are overall relatively low. Without this minimum, weekly and daily movements in implied ratings may be volatile, as small movements in EDF levels cross implied rating mapping boundaries. Since we build the implied ratings mappings out from Baa, the median EDF value for this rating category is never adjusted.

Finally, geometric means (of the adjusted median EDF values of neighboring rating categories) are employed once more – this time, to determine the upper and lower bounds for each implied rating. EDF-implied rating tables updated monthly using spot medians reflect dynamic mappings. Also available on CreditEdge are static mappings based on long-run default rates.

EDF-implied ratings are assigned by finding the appropriate implied rating table EDF range.

For More Information

To learn more about the EDF model and its applications, please contact our experts at clientservices@moodys.com.