

MODELING METHODOLOGY

FROM MOODY'S ANALYTICS
QUANTITATIVE RESEARCH

Authors

Christopher Crossen
Shisheng Qu
Xu Zhang

Contact Us

Americas
+1-212-553-1653
clientservices@moodys.com

Europe
+44.20.7772.5454
clientservices.emea@moodys.com

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

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

Validating the Public EDF Model for North American Corporate Firms

Abstract

In this paper, we validate the performance of the Moody's Analytics Public EDF™ (Expected Default Frequency) model for North American corporate firms during the last decade, including the recent credit crisis and its recovery period. We divide the decade into two sub-periods: an early period, 2001–2007, and a later one, 2008–2010, and then compare the model's performance during these two periods. We focus on the model's ability to differentiate between good and bad firms, its comparison to agency credit ratings, the timeliness of its default prediction, and its accuracy of levels.

Overall, the EDF model's predictive power during the recent sample period has been consistent with its previous longer history, outperforming alternative risk measures, including agency credit ratings. On average, the model provides an early warning signal at least 12 months before default occurs. EDF levels were conservative (higher than subsequently realized default rate) before the crisis when compared with later-realized default rates, and levels were statistically consistent with later-realized default rates.

We find that EDF credit measures perform consistently well across different time horizons. Our tests indicate that EDF credit measures provide a very useful measure of credit risk that can be applied throughout North America.

Table of Contents

1	Overview	5
2	The Predictive Power of EDF Credit Measures	6
	2.1 Data	6
	2.2 Validation Results	7
	2.3 Rated Firms	8
3	EDF Credit Measures as Early Warning Signals.....	10
4	Level Validation	12
	4.1 Comparing EDF Levels and Realized Default Rates for EDF Groups	12
	4.2 Validating EDF Levels with Simulated Levels.....	14
5	Conclusion	15
	References.....	17

1 Overview

In this paper, we present the results of a study validating the performance of Moody's Analytics' EDF credit measures for North American corporate firms¹ during the last decade, including the recent credit crisis and subsequent recovery. We divide the decade into two sub-periods: an early period, 2001–2007, and a later one, 2008–2010, and then compare the model's performance during these two time spans. We focus on the model's ability to differentiate between good and bad firms, its comparison with agency credit ratings, the timeliness of its default prediction, and its accuracy of levels.

We validate the EDF model on a regular basis. Our 2007 study examined the period from 2001 through 2006 (Dwyer and Korablev, 2007); our 2009 study focused on the model's performance during the peak of the recent credit crisis (Korablev and Qu, 2009).² This paper details our latest study, which extends our review of the model's performance into the "recovery" period of the credit crisis, including data through 2010. Credit analysis becomes more difficult during high default rate periods, but the severity and complexity of this recent downturn makes default prediction more challenging. In this study, we follow our existing model performance testing methodology, as described in Bohn, Arora, and Korablev (2005) and Dwyer and Korablev (2007).

Overall, we find the EDF model's predictive power during the recent sample period has been consistent with the previous longer history, outperforming alternative risk measures, including agency credit ratings. On average, the model provides an early warning signal at least 12 months before default occurs. EDF levels were conservative (i.e., not too low) before the crisis compared with later-realized default rates, and levels were statistically consistent with subsequently realized default rates.

This paper is organized in the following way:

Section 2 presents results for the EDF model's predictive power for North American corporate firms.

Section 3 discusses EDF credit measures as early warning signals for default.

Section 4 shows the level validation results.

Section 5 provides concluding remarks.

¹ "North American corporate firms" excludes financial firms.

² See Korablev and Qu (2009), "Validating the Public EDF Model Performance During the Recent Crisis."

2 The Predictive Power of EDF Credit Measures

One of the most important applications of a default prediction model is to differentiate bad firms (i.e., genuinely distressed) from good firms. In this paper, to test the power of the EDF model, we use a well-known approach: the Cumulative Accuracy Profile (CAP). This approach is summarized by a measure known as the Accuracy Ratio (AR). Typically, the higher the AR, the better the model performs. In extreme cases, for a totally random model that bears no information on impending defaults, $AR=0$. For a perfect model, $AR=100\%$.³

2.1 Data

In this study, to assess the predictive power of the EDF credit measure, we compute Accuracy Ratios on the North American corporate sector. We compare Accuracy Ratios for the EDF model in predicting defaults between 2001–2007 and between 2008–2010. All Accuracy Ratios are for a one-year horizon. Our findings show that the inclusion of the crisis period does not reduce the predictive power of the EDF credit measure.

In all tests, we use defaults included in the Moody's Analytics default database, collected and updated daily from numerous printed and online sources worldwide.⁴ As a result, Moody's Analytics employs the most extensive public company default database available. Nevertheless, small public companies often disappear without news or record before they default, or they do not publicly disclose missed payments, which creates a number of hidden defaults, a common challenge often faced by default risk researchers. To reduce this problem of hidden defaults, in many of our tests, we restrict the sample to firms above a certain size threshold, where we believe hidden defaults are less of an issue. We typically use a size threshold of greater than \$30 million or \$300 million in annual sales, depending upon the context.⁵

In this study, unless otherwise indicated, we calculate Accuracy Ratios by pooling all firm-month observations before ranking them, so that the same EDF value at different time points is assigned the same risk ranking. We track each firm-month EDF observation for 12 months and assign a flag of default or survival. The Accuracy Ratio is computed using the EDF value and default flag pairs, ignoring when the EDF is measured.

The majority of the North American corporate sector consists of U.S. firms. However, we also collect data on firms outside the U.S., including Canada, Bermuda, the Cayman Islands, the Bahamas, Panama, British Virgin Islands, Netherlands Antilles, and Belize. During 2001–2010, there were 1,478 unique default events, with 1,219 during 2001–2007 and 259 during 2008–2010. Table 1 shows the countries and the number of firm-months in each country that constituted the North American module in Moody's Analytics Credit Monitor® and Moody's Analytics CreditEdge® between 2001 and 2010. As the table shows, outside the U.S., Canada has the largest number observations in the sample, followed by Bermuda and the Cayman Islands.

³ For more details on Accuracy Ratios and a related measure of Receiver Operating Characteristic (ROC), please see Keenan and Sobehart (2000) and Dwyer and Korablev (2007).

⁴ We utilize government filings, government agency sources, company announcements, news services, specialized default news sources, as well as sources within financial institutions to ensure, to the greatest extent possible, that we account for all defaults.

⁵ All figures are in U.S. dollars, unless noted otherwise. We measure size by total annual sales for North American non-financial firms. Where the firm's total sales number is not available, we use book assets.

Table 1 Countries in the North American Database: 2001–2010

Country	Number of Observations (firm-months)
United States	376,100
Canada	57,315
Bermuda	2,262
Cayman Islands	649
Bahamas	370
British Virgin Islands	269
Netherlands Antilles	235
Panama	173
Belize	70

2.2 Validation Results

We track the solvency of each EDF observation for 12 months. During the 2001–2007 period, we use EDF credit measures between December 2000 and December 2006 and default events between January 2001 and December 2007. Similarly, for the 2008–2010 period, we use EDF credit measures between December 2007 and December 2009 and default events between January 2008 and December 2010. We also compute Altman’s Z-Scores for these periods, using the same time frames as with the EDF measures. In Figure 1, the left panel shows the EDF model’s performance during 2001–2007. The right panel presents model performance during 2008–2010. Both samples include only firms with greater than \$30 million in annual sales. During the later, crisis period, the EDF credit measure’s Accuracy Ratio is 81.9%, compared with the 83.8% Accuracy Ratio during the 2001–2007 period.⁶ Although problems may arise when comparing Accuracy Ratios from two different data samples, comparing the recent sample, which includes the recent credit crisis, with the prior sample, suggests that the EDF model performs consistently over time and in different credit cycles.⁷ The EDF credit measure also performs substantially better than the Z-Score during both periods.

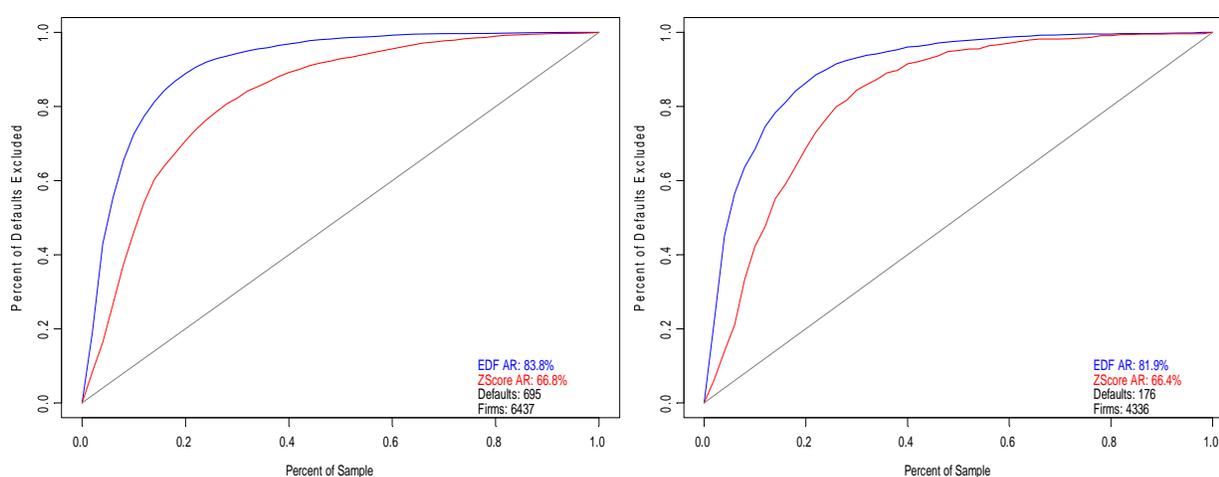


Figure 1 CAP Curves Comparing EDF Credit Measure Historical Performance for North American Corporate Firms, Annual Sales > \$30 million: 2001–2007 versus 2008–2010.

⁷ As a result of the entry and exit of firms over time, the sample of firms used for the validation changes over time to some extent also.

Table 2 shows the Accuracy Ratios for the EDF credit measure and the Z-Score at the end of each year between 2000 and 2009. We can see that the EDF credit measure outperforms the Z-Score by large margins during all years.

Table 2 Accuracy Ratios for EDF credit measures and Z-Scores by year/month

Date (end of)	Numbers of Companies	Numbers of Defaults	Default Rates	Accuracy Ratios EDF	Accuracy Ratios z-score	Accuracy Ratio Differences
2000/12	4,653	253	5.4%	72.0%	66.0%	6.0%
2001/12	4,210	157	3.7%	78.7%	67.6%	11.2%
2002/12	3,838	99	2.6%	82.3%	64.5%	17.8%
2003/12	3,810	51	1.3%	86.2%	71.1%	15.0%
2004/12	3,857	49	1.3%	81.5%	59.9%	21.5%
2005/12	3,878	23	0.6%	85.5%	70.1%	15.4%
2006/12	3,919	21	0.5%	93.9%	60.7%	33.2%
2007/12	3,836	57	1.5%	82.9%	75.0%	7.9%
2008/12	3,726	88	2.4%	86.7%	63.8%	22.9%
2009/12	3,572	24	0.7%	87.6%	62.1%	25.5%

2.3 Rated Firms

In this section, we focus on the subset of firms with both EDF credit measures and Moody's Investors Service credit ratings. Figure 2 presents CAP Curves. During 2001–2007, shown in the left panel, the Accuracy Ratios of the EDF model and credit ratings were 89.6% and 76.5%, respectively, derived from 196 defaults and 1,475 unique firms. During 2008–2010, shown in the right panel, the Accuracy Ratios were 84.5% and 84.4%, respectively, for the EDF model and credit ratings, derived from 66 defaults and 965 unique firms.⁸ Although we see a cross-over in the power curves for EDF credit measures and ratings during the recent crisis period, EDF credit measures still perform well month-by-month.

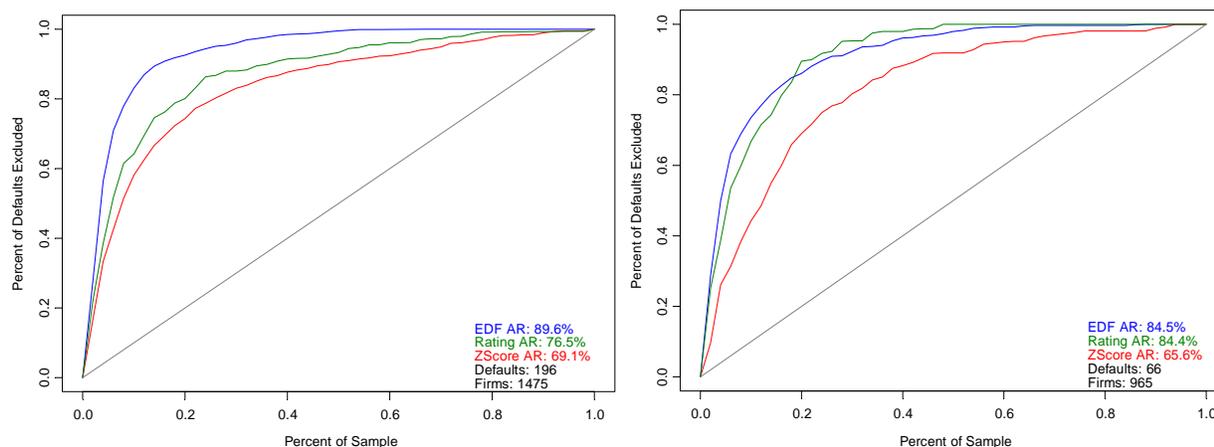


Figure 2 CAP Curves Comparing EDF Credit Measure Historical Performance for North American Corporate Firms Rated by Moody's Investor Services: 2001–2007 versus 2008–2010.

Figure 2 also shows that the advantage EDF measures have over agency ratings is much smaller during the crisis period when compared to the longer history. This finding is mainly due to the fact that the EDF credit measure is a more

⁸ Ratings may be withdrawn before a default event occurs. We count a firm as a rated default if it was rated within 12 months prior to its default.

dynamic measure of current credit quality, affected more by a turbulent economic environment. In measuring Accuracy Ratios, we rank companies only by their EDF values and do not consider their credit environments. Therefore, for example, a company with a 2% EDF during mid-2007 is treated as risky as a different company with a 2% EDF during early 2009, despite the former being equivalent to a Caa3 company and the latter a Ba3 company at their respective points in time. In reality, when rank-ordering companies, most people compare risks at the same point in time.

To capture predictive power in this relative sense, we also compute Accuracy Ratios differently: we rank EDF measures by observation time, pool the outcome rankings across time, and compute the ARs. The same is applied to ratings, but as ratings are relatively static, their risk rankings do not materially change. Figure 3 shows results, where the advantage of the EDF credit measure is more visible, and the cross-over in Figure 2 is no longer present.

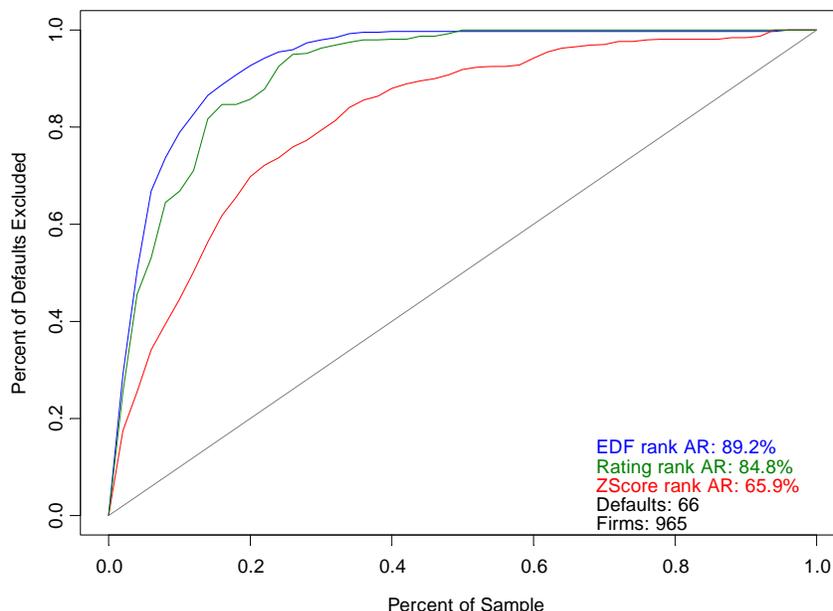


Figure 3 EDF Credit Measures Ranked by Observation Time, with Rankings Pooled Across Time, for North American Corporate Firms Rated by Moody's Investor Services 2008–2010.

Looking at the same issue from a different angle, Table 3 shows Accuracy Ratios computed from only cross-sectional data, at the end of each year. While we see that EDF credit measures outperform ratings in all of these cross-sections, the outperformance tends to be smaller prior to large unexpected shocks to the market. For example, at the end of 2007, the market had experienced a five year increase, but the following year marked a significant decline. At the end of 2009, the market began recovering from a bottom point, but with great reluctance, and the EDF values turned out to be not optimistic enough relative to quickly falling default rates. Agency ratings, being relative measures and much more static, are much less affected by such large, systematic shocks.

Table 3 Accuracy Ratios for EDF credit measures and Moody's Investor Service credit ratings by year/month

Date (end of)	Numbers of Companies	Numbers of Defaults	Default Rates	Accuracy Ratios		Accuracy Ratio Differences
				EDF	Rating	
2000/12	1,097	76	6.9%	74.5%	66.9%	7.6%
2001/12	1,015	46	4.5%	83.5%	64.1%	19.4%
2002/12	969	34	3.5%	90.6%	85.4%	5.2%
2003/12	927	15	1.6%	85.3%	84.4%	0.9%
2004/12	898	15	1.7%	91.4%	77.4%	14.1%
2005/12	935	6	0.6%	91.5%	82.4%	9.1%
2006/12	890	3	0.3%	97.7%	94.1%	3.6%
2007/12	872	18	2.1%	84.8%	82.8%	2.0%
2008/12	837	40	4.8%	90.6%	84.5%	6.1%
2009/12	816	7	0.9%	92.4%	91.4%	1.0%

3 EDF Credit Measures as Early Warning Signals

To test the timeliness of default prediction as an early warning signal, we create a sample of defaulted firms from January 2008 through December 2010. We compute the 25th, 50th, and 75th percentiles of the EDF credit measure for these defaulted firms dating back to December 2001. We also compute the same percentiles for the entire sample, and then plot two sets of percentiles on the same graph. If the EDF credit measure provides early warning signals, we expect the default sample EDF credit measure distribution to move higher and away from that of the entire population as they approach default dates. This finding is indeed what we observe.

In the North American corporate sector, 172 firms defaulted between January 2008 and December 2010. Figure 4 presents EDF credit measure percentiles for these defaulted firms. The red lines represent the 25th, 50th, and 75th percentiles of EDF credit measures for companies that defaulted between January 2008 and December 2010. The blue lines represent percentiles for the entire sector.

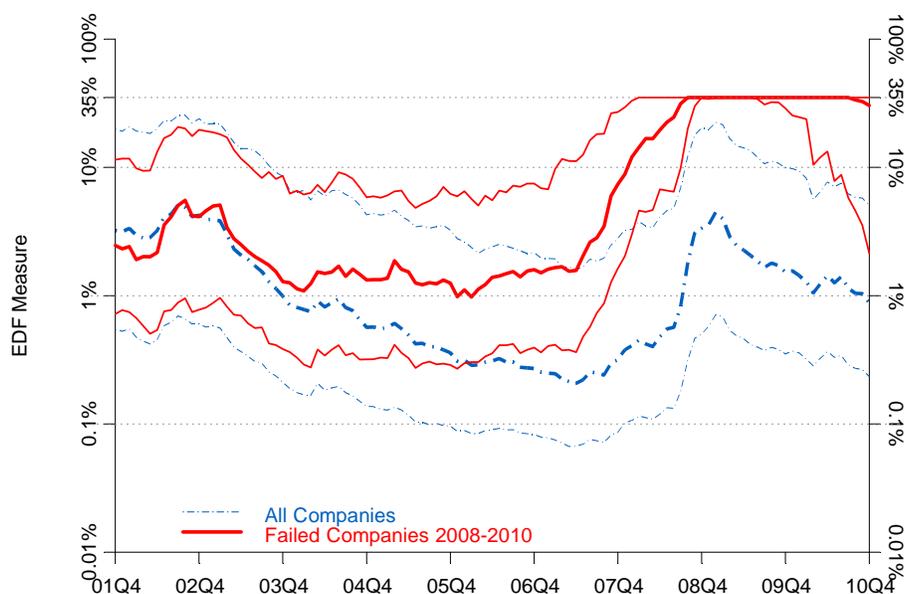


Figure 4 Early Warning and Distribution of EDF Credit Measures for North American Corporate Firms: 25th, 50th, and 75th percentiles.

These two distributions are distinctly different. As shown in Figure 4, beginning in mid-2004, defaulters were riskier than the rest of the sample and usually had higher EDF credit measures than non-defaulters, well before default. As the entire sector improved between the end of 2002 and mid-2007, defaulters actually began to deteriorate during mid-2006. The top 25th percentile of the defaulters' EDF measures began increasing as far back as late-2005. The speed of the deterioration increased in mid-2007. Defaults were realized between January 2008 and December 2010. The entire sector's risk began to increase beginning in mid-2007 and began to decrease beginning in early 2009.

In Figure 4, we also observe a steep drop in the 25th percentile of the defaulter's EDF credit measures, beginning in early 2010. This finding is mainly due to some surviving companies emerging from bankruptcy that were then relisted in the equity market, and thus, their EDF levels improved significantly. For example, Visteon Corporation's EDF credit measure changed from 35% in August 2010 to 0.56% in December 2010, a result of its re-organization and the relisting of its equity in October 2010. Another example is GSI Group, Inc., which filed bankruptcy in November 2009, entered reorganization, and emerged from bankruptcy in May 2010. Its EDF credit measure dropped from 21.7% in August 2010 to 0.89% in December 2010.

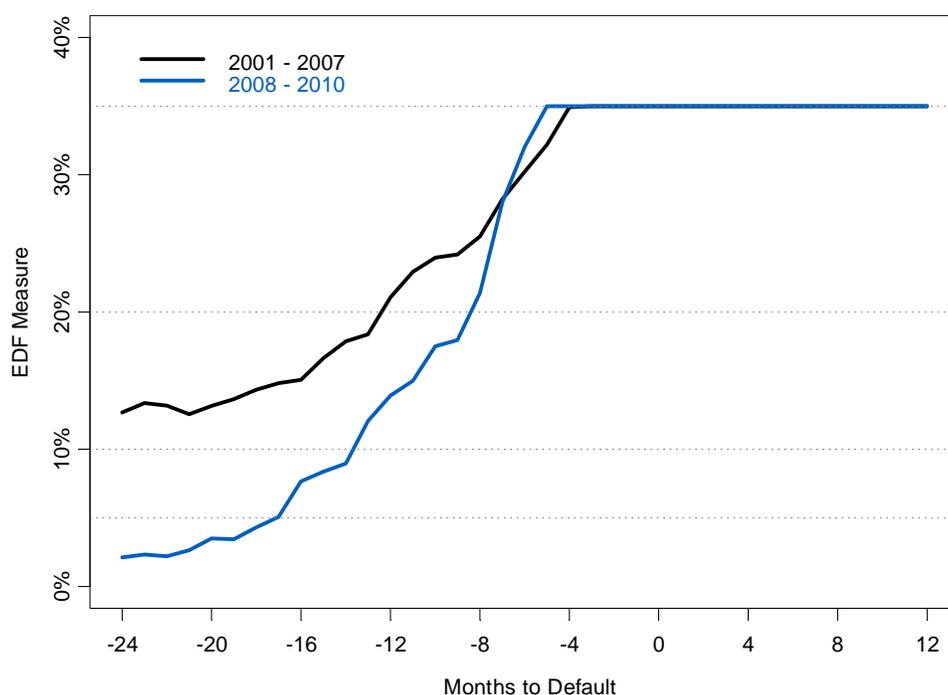


Figure 5 Median EDF Credit Measures for North American Corporate Defaulters.

Another way to test the timeliness of default prediction is to measure the number of months before an impending credit event that the EDF measure provides a signal of deteriorating credit quality. To test timeliness, we create a sample of defaulted firms, retaining monthly observations from 24 months prior to default, and then compute the median EDF credit measure by months to default. In Figure 5, we overlay and compare the median EDF credit measures of the two periods (2001–2007 and 2008–2010).

Figure 5 demonstrates that, in the event of default, the early warning performance of EDF credit measures, including the recent recession, is similar to historical performance. For both periods, the EDF credit measure is elevated more than 12 months prior to the credit event and continues increasing steadily. Figure 5 shows that the median EDF level 24 months prior to default is lower for the 2008–2010 period compared with the 2001–2007 period. This difference is not surprising, given that the two years prior to the crisis, 2005–2007, were a relatively benign period. Additionally, in 2008–2010, the slope of the increase in median EDF level steepens as early as 18 months before default occurs.

4 Level Validation

Level validation measures how well the EDF model's predicted default rates track realized default rates. EDF values declined consistently since 2002, reaching their lowest levels in mid-2007. The subsequent credit crunch saw a dramatic turn in the credit environment. After observing an unprecedented increase in EDF credit measure levels and elevated default rates from 2007 through 2009, we ask the following question: did the EDF model underestimate risk during more benign times? In this section, we provide evidence showing that EDF levels were not too low relative to the level of defaults observed later.

We use the following distinct approaches to assess EDF levels against observed default rates:

- We group firms by their EDF levels and compare realized default rates for each group. This grouping can be done at any given point in time and can also be aggregated across time. Either way, we find that the EDF credit measure remains consistently conservative (i.e. high relative to realized defaults).
- Even if a PD model truly describes default risk, because defaults are random and correlated, realized default rates for a portfolio can differ from the average PD predicted by this "true model." We simulate defaults using EDF values and asset correlation and compare the simulated default rate distribution with the observed default rate. We find that we cannot reject the hypothesis that the EDF model is the true default probability model.

4.1 Comparing EDF Levels and Realized Default Rates for EDF Groups

We begin by comparing the one-year EDF values with realized default rates for firms grouped by their EDF levels. The one-year EDF credit measure at time t is designed to describe the expected default frequency for the 12 months following time t . If the model is correct, for a group of firms with a given EDF level, the realized default rate from the group during the following 12 months should be near the average EDF level. If the model underestimates risk, the realized default rate should be higher than the average EDF level.

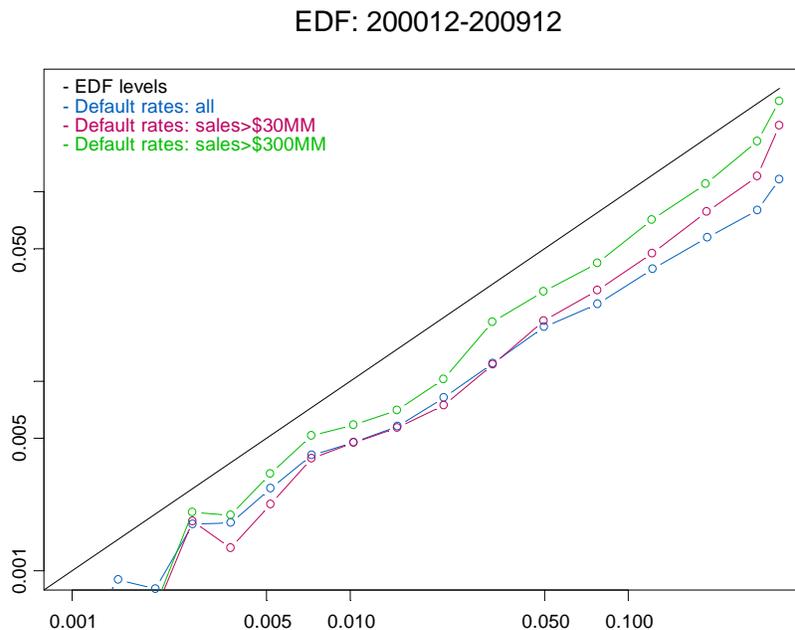


Figure 6 Distribution of Average EDF Levels and Default Rates for North American Corporate Firms: 2001–2010.

Figure 6 provides a comparison of North American corporate firms. As shown in the figure, traditionally, EDF measures have been consistent with realized default rates for North American firms. We take EDF credit measures between January 2001 and December 2009 and defaults between January 2001 and December 2010, allowing the last EDF credit measure 12 months to default. The panel data is grouped into 20 EDF groups with equal sample sizes (5% each). Figure 6 indicates that EDF levels are higher than observed default rates (even at different annual sales restrictions) as expected, due to the hidden defaults issue.

The hidden defaults issue refers to the failure of a default dataset to capture all economic defaults. This failure can occur for various reasons. For example, when a debt extension occurs, it is difficult for an outsider to know if the extension is caused by the borrower's inability to pay or by legitimate business need. In other cases, when the loan amount is small, failure to pay is simply written off by the bank, and no public announcement is released. When default data collection relies on public information to identify defaults, many default events may go missing. This is particularly true for smaller firm borrowers that draw little public attention.⁹

Moody's Analytics' team of specialists aggregates default data utilizing multiple information sources including, but not limited to: bankruptcy newsletters, rating agency debt monitoring publications, news media and news search engines, corporate regulatory filings, internet browsing, and targeted searching. As of January 2011, the database has recorded more than 8,000 unique default events for public firms or firms that were public before default occurred. We collected many of these defaults in real-time during the past two decades. Despite being the largest public default database we are aware of, we suspect there remain a significant number of defaults not captured in the data, because in many cases, distressed borrowers work out deals privately with lenders, drawing little attention in the media or by data collectors. As there is generally less information available for smaller companies, we believe the hidden default problem is larger for smaller companies. In calibrating the EDF model, we consciously employ only large companies in mapping distance-to-default to EDF levels in order to circumvent this problem. In addition, we construct the mapping so the EDF measure is conservative relative to the long-term average default rate, even in the large company sample.

The sample underlying Figure 6 constitutes all North American corporate firms during 2001–2010. Overall, we see that EDF levels are higher than realized default rates, and, as we limit the sample to larger and larger companies, the realized default rates at given EDF levels, increase as well, indicating that limiting the sample to larger firms reduces the risk of hidden defaults. This long-term picture provides a benchmark for a more recent snapshot reported in Figure 7.

Figure 7 shows the more interesting, recent period, 2008–2010. During this time, EDF levels increased significantly from historic lows in the wake of the most severe economic turmoil since the Great Depression. In addition, corporate defaults increased. There were concerns that EDF levels might be too low overall; Figure 7 shows that this was not the case.

In both Figure 6 and Figure 7, for the riskiest group (Annual Sales > \$300 million), where most defaults occurred, the EDF levels and default rates are approximately the same on the large company sample.

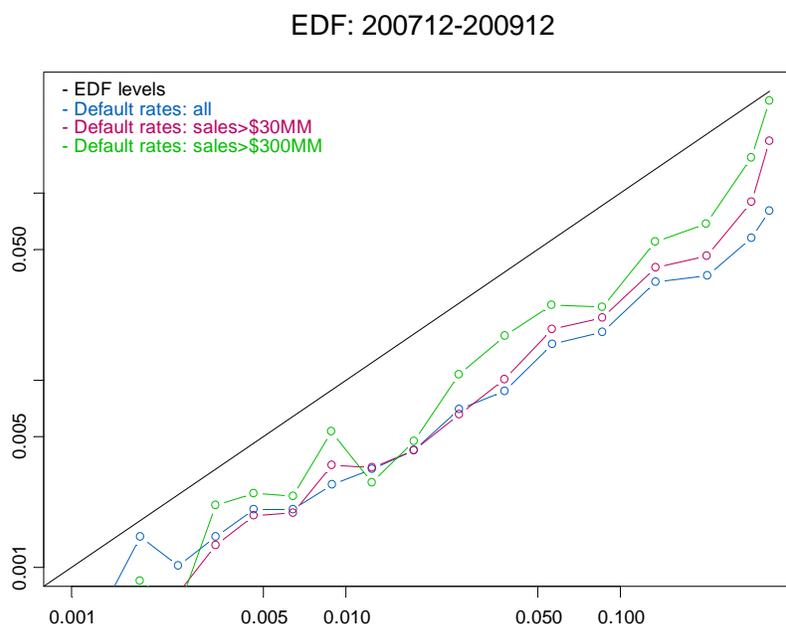


Figure 7 Distribution of Average EDF Levels and Default Rates for North American Corporate Firms: 2008–2010.

⁹ See Stein and Dwyer (2005) and Dwyer and Qu (2007) for more information regarding hidden defaults.

Given the overall higher EDF levels than realized default rates, it is fair to ask a basic question: are EDF credit measures too conservative? The answer is no.

Figure 8 shows the realized EDF levels compared with default rates for default events during 2001 and 2008. We can see that “over-estimation” is much less of a problem in 2008, as defaults increased. While overall, EDF credit measures are conservative, they are not too conservative, in the sense that, given the distance-to-default level, such high EDF levels could happen, and they actually did occur.

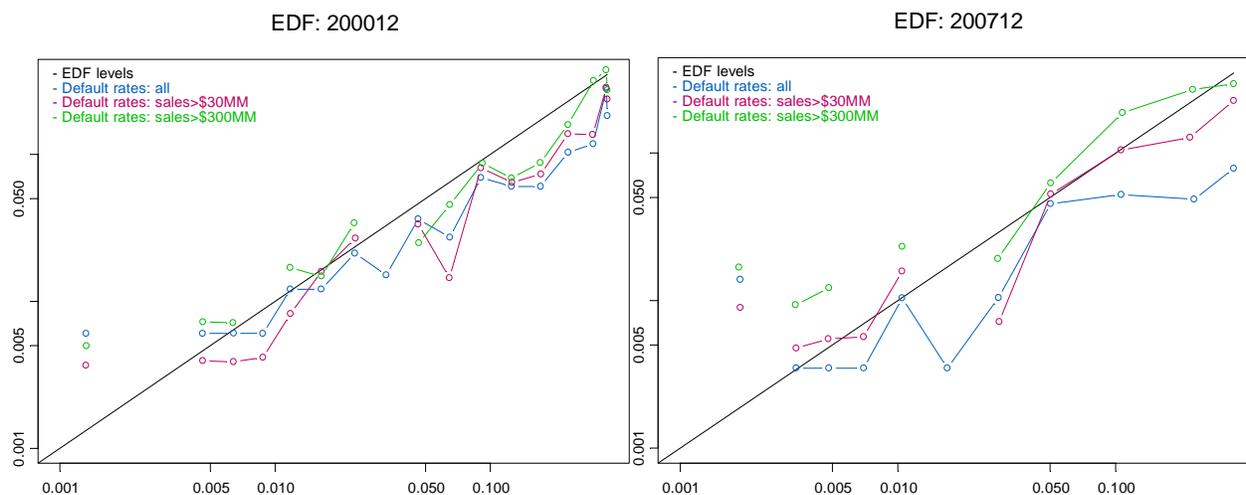


Figure 8 Distribution of Average EDF Levels and Default Rates for North American Corporate Firms: 2000/12 and 2007/12.

4.2 Validating EDF Levels with Simulated Levels

In this section, we use a structured approach to test whether or not we can statistically reject the hypothesis that the EDF model is a true model for default prediction.

EDF credit measures are probability measures, and, for a given portfolio of companies, the model implies a distribution of possible default rates. We test whether or not the observed default rate is consistent with such a default rate distribution.

We use a simple numerical example to illustrate the nature of the test. Suppose we have 100 independent companies, each with a default probability of 10%. The expected default rate is 10%, and it is possible, but not likely, to have a default rate of less than 5% (the p-value is only 5.75%).¹⁰ When we actually observe a 5% default rate, we would reject, at a 10% confidence interval, the hypothesis that the default probabilities are 10%.

When companies are correlated, the implied default rate distribution is wider, in that, the likelihood is higher for the realized default rate to deviate from the mean prediction. For example, for two independent firms, each with a 50% true default probability, the likelihood of observing no default, a negative deviation from the expectation of 50%, is only 25%. However, if these two firms are perfectly correlated, the likelihood of no default increases to 50%.

Based on the above concepts, Kurbat and Korabiev (2002) developed a method that uses realized defaults for testing default probability models. Specifically, they assume firm values are correlated with one common factor, and obtain a default rate distribution by simulating random realizations of a common factor and firm-specific factors. In simulations, each firm’s unconditional default probability is kept at the level predicted by the subject probability of default (PD) model. In this study, we use the same approach to test if the observed default rate is consistent with the default rate distribution implied by the EDF credit measures. For each year in the sample, we take EDF values calculated at the

¹⁰ The likelihood of less than n% default is $\sum_{i=0}^n \frac{100!}{i!(100-i)!} (0.1)^i (0.9)^{100-i}$.

beginning of the year, use them to simulate a distribution of default rates, and compare the actual default rate during the year against the resulting distribution. We use a pair-wise asset correlation of 0.19, calibrated using long-term data from the Moody's Analytics Global Correlation Model (GCorr). For each year, we run 1,000 simulations to create a distribution of 1,000 simulated default rates.

Figure 9 presents results. To avoid the hidden defaults issue, we limit our sample to companies with annual sales greater than \$300 million. As shown in Figure 9, the realized default rate was mostly lower than the median EDF level during the past ten years, but it remained within the intervals bounded by the 10th and 90th percentiles of the simulations, with the exception of 2010.

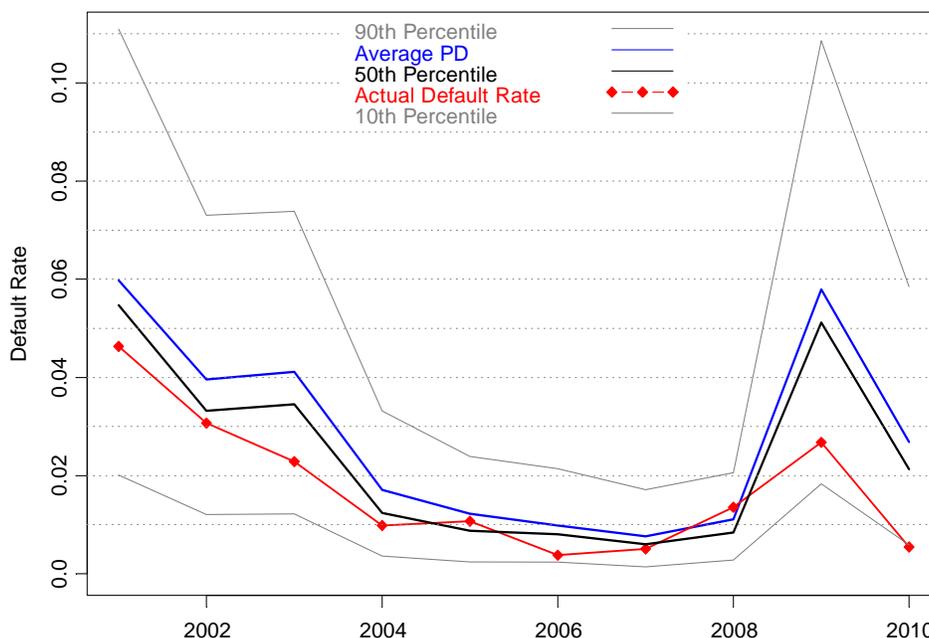


Figure 9 Default Rate for North American Corporate Firms (Annual Sales > \$300 million), 2001-2010.

We can conclude from these results that the observed default rates are consistent with EDF model predictions for all years.

5 Conclusion

In this study, we test the public EDF model for North American corporate firms using three major performance measures: Accuracy Ratios in default prediction, early warning signals, and default risk levels, with special focus on the most recent credit crisis and recovery period.

In Accuracy Ratio testing, we find that EDF credit measures are as powerful as they have been historically in their ability to discriminate good firms from bad firms over time, and that they consistently outperform agency credit ratings.

We find that the EDF model provides ample early warning signals. The distribution of EDF levels for defaulters begins to emerge out of the entire population distribution more than 12 months before defaults occur.

EDF levels were conservative (i.e., not too low) before the crisis, compared with later realized default rates, and levels were statistically consistent with later realized default rates. We find that the EDF levels are consistently higher than observed default rates, due to downwardly-biased default rate observations, caused by the hidden defaults issue, as well as conservatism built into the EDF model. Over the longer history, the realized default rate, with better default coverage, lies within the prediction interval, and we cannot reject the hypothesis that the EDF values are true measures of default risk.

Acknowledgements

We are grateful to everyone who contributed to this paper. All remaining errors are, of course, our own.

Originally published May 2011, revised December 2011.

Copyright © 2011 Moody's Analytics, Inc. and/or its licensors and affiliates. All rights reserved.

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, 2005a.
- Crosbie, Peter and Jeffrey Bohn, "Modeling Default Risk." Moody's KMV, Revised December 2003.
- Dwyer, Douglas and Irina Korablev, "Power and Level Validation of Moody's KMV EDF™ Credit Measures in North America, Europe, and Asia." Moody's KMV, September 2007.
- Dwyer, Douglas and Shisheng Qu, "EDF™ 8.0 Model Enhancements." Moody's KMV, January 2007.
- Keenan, Sean and Jorge Sobehart, "A Credit Risk Catwalk." Risk, July 2000.
- Korablev, Irina and Shisheng Qu, "Validating the Public EDF Model Performance During the Recent Credit Crisis," June 2009.
- Kurbat, Matt and Irina Korablev, "Methodology for Testing the Level of the EDF Credit Measure." Moody's KMV, August 2002.
- Stein, Roger and Douglas Dwyer, "Inferring the Default Rate in a Population by Comparing Two Incomplete Default Databases." Moody's KMV, October 2005.