MOODY'S ANALYTICS

VIEWPOINTS

Banks and their EDF Measures Now and Through the Credit Crisis: Too High, Too Low, or Just About Right?

Capital Markets Research, Inc.

Authors

David W. Munves, CFA
1.212.553.2844
david.munves@moodys.com

Allerton (Tony) Smith
1.212.553.4058
allerton.smith@moodys.com

David T. Hamilton, PhD
1.212.553.1695
david.hamilton@moodys.com

Research Assistance

Ervis Deda
1.212.553.1404
ervis.deda@moodys.com

Lee Chua
1.212.553.4383
lee,chua@mkmv.com

Summary

Financial institutions, particularly banks, were at the heart of the credit crisis and subsequent recession, and defaulted at unprecedented rates. It will be a long time before names like Lehman Brothers, Bear Stearns, and Northern Rock fade from the memories of investors and risk managers. Not surprisingly, the experience has redoubled interest in finding effective and efficient ways to provide early warning of credit distress for such entities.

Moody's Analytics' Expected Default Frequency (EDF™) metrics are one of the most widely used types of probability of default measures in quantitative credit risk analysis. In this paper we explore the performance of the EDF model for public firms as it relates to banks. The time is right: the credit crisis and subsequent recession have only recently ended, leaving us with a mass of data and experiences to analyze. Our key findings are as follows:

» EDF credit measures did a good job in rank ordering defaulters during the crisis: financial institutions that subsequently defaulted had high EDF measures relative to those of their peers.

» However, the EDF measures for many financial institutions that defaulted were low in absolute terms, leading to the impression in some quarters that EDF metrics had not performed as expected. The data shows that this was not the case — the default rate for financials was in line with the level of risk indicated by the model. Users of the public firm EDF model can improve their surveillance of entities with low nominal EDF measures by comparing the movements of individual entities' EDFs to those of their sectors.

» EDF metrics for many banks remain well above pre-crisis levels. To a large degree this is due to the fact that risk levels are indeed high. However, we find that elevated measures also reflect the model's focus on risk across all levels of a firm's capital structure, as well as other factors particular to banks. One strategy is to focus on changes in EDF measures, rather than on absolute levels. Another is to reclassify some types of liabilities and recalculate the EDF metric based on the adjusted inputs.

EDFs for Financial Institution: Defaulters vs. the Sector

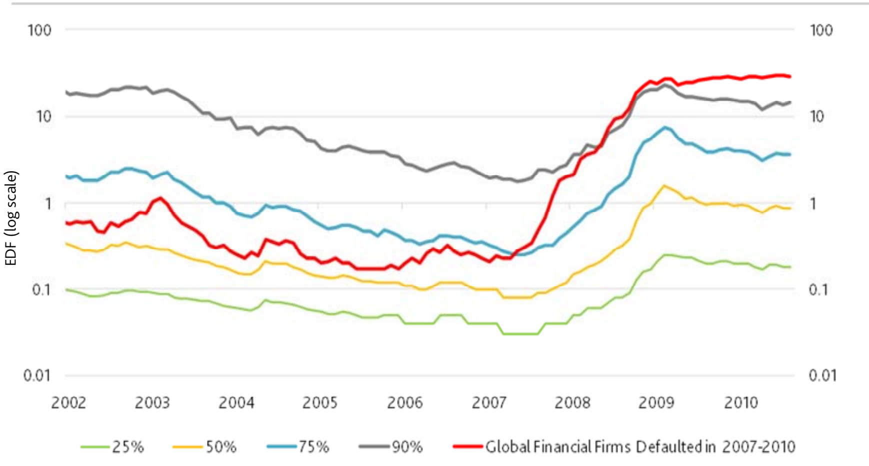

Table of Contents

Introduction

"How can an institution like Deutsche Bank have a B2-equivalent EDF?" is a frequently heard objection

Challenges in interpreting EDF credit measures for banks

Two characteristics of EDF measures for banks often come up in conversations with market participants. One is that they were too low heading into the recession. The other is that having come through the worst financial crisis in 80 years, they're now too high: "how can an institution like Deutsche Bank have a B2-equivalent EDF?" is a frequently heard objection.[1] Implicit in such comments is the general impression that models like the one behind public firm EDFs "don't work" for banks. This stems from considerations such as the risk of abrupt default due to the loss of market access, or the different liability structures of banks (as opposed to corporates).

This paper tackles such issues head-on. We start with a review of the model's performance for financial institutions before and during the recession, including a case study of Lehman Brothers. We then analyze the reasons why many financial institutions' EDF measures remain so elevated.

We also include "users' tips" for employing EDF measures for banks, which take into account some of the aforementioned factors. Such ideas are a major part of our discussions with clients. They reflect how EDFs and other quantitative metrics are used in many risk management processes — as tools to select entities on which to focus, and as inputs into credit decisions. That is, EDFs are means to an end, with the end being more efficient and better risk management.

For some readers this will be their first encounter with EDF measures, so a brief explanation is in order at this point.

Expected Default Frequency metrics are Moody's brand name for probability of default estimates. This paper is about the *Public Firm model*, i.e., the model covering firms with traded equity and public financial statements. Public Firm EDF measures are based on information from firms' capital structures and equity prices. They were first produced in the early 1990s by KMV. KMV was purchased by Moody's Corporation in 2002 and subsequently renamed Moody's-KMV. There is a separate Private Firm EDF model. As the name suggests, it produces EDFs for firms without publicly traded shares.[2] Both types of EDF metrics are estimates of "physical" PDs, as distinct from "risk-neutral" PDs. The latter are useful in asset pricing but not risk management, because they include the effect of the market price of risk. As a result, risk-neutral PDs are always higher than physical PDs and therefore overstate real default risks.[3] EDF credit measures represent pure default risk, with no consideration of loss-given default rates. The remainder of this paper assumes a general understanding of Public Firm EDF metrics and how they're calculated. Readers in need of a refresher should refer to Appendix 1.

Focus on banks

Financial institutions are a varied group, encompassing banks, insurance companies, finance companies, and other sub-groups. Banks make up 34% of financial institutions with EDF measures. Outright defaults of large banks have been rare, due in part to government interventions. However, when they do occur the results can be near cataclysmic. Banks also remain much in the news, and their central role in the global economic system is unchallenged. So while we assess EDF performance across all financial institutions, reflecting the need for a large number of observations for such studies, this paper is mostly about banks.

[1] This and other company-specific EDF levels reflect data as of November 29, 2010. The B2 equivalent EDF is calculated from a spot mapping of EDFs to Moody's ratings for a broad population of firms. Footnote 11 provides more details about this.

[2] For details about the EDF private firm model please see Dwyer and Kocagil (2004).

[3] For details about physical vs. risk-neutral PDs, please see Dwyer *et al* (2010)

Simply put, the passage of time should reveal that the defaulters were concentrated among the entities that had relatively high EDF measures some time prior to their defaulting

EDF Metric Performance Update

A lot of the research around the public firm EDF model consists of analyzing its performance based on various criteria. Thus, this section largely represents an update of existing papers,[4] which cover EDF performance through the end of 2008.

A key goal: rank ordering of default risk

A key test of a default risk model is its ability to rank order entities by their risk scores, in this case their EDF measures. Simply put, the passage of time should reveal that the defaulters were concentrated among the entities that had relatively high EDF measures some time prior to their defaulting. This is also referred to as a test of the model's statistical power. A second test is whether EDF levels are in line with subsequent default rates. We assess prospective default risk identification ability by analyzing aggregated sets of data to ascertain statistical performance. Doing so avoids the (sometimes temping) alternative of seizing on single examples to support an argument.

CAP curves and accuracy ratios

A common measure of model power in rank ordering default risk is the cumulative accuracy profile, or CAP, curve. Figures 1a and 1b show CAP curves for the EDF financial institutions data set, with the former covering the pre-crisis period (through 2006) and the latter showing subsequent events. The horizontal axes of the CAP curves rank-order the population being analyzed by their EDF measures at the beginning of a 1-year period. The ranking from the highest EDF metrics (on the left) to the lowest (on the right), and includes both defaulters and non-defaulters.[5] The measurement units correspond to the percentile ranking of the firms based on their EDF metrics. The vertical axes show the cumulative proportion of firms that defaulted during a 12-month period, *based on their EDF rankings (in percentiles, among the entire population) at the beginning of the period.* Crucially, this builds in the notion of early warning, since the defaulters' relative rankings are not based on their EDFs measures just prior to their declarations of bankruptcy.[6] Thus, in Figure 1a we see that the 10% of the population with the highest EDF metrics experienced 84% of the defaulters, the 20% of the group with the highest EDF measures encompasses 91% of the defaulters, and so on. The two figures are not strictly comparable, since the data sets differ (some institutions only existed in the first period, and some only in the second). But they're the same *types* of entities, so we can safely compare them to each other.

Figure 1a and 1b — Power Curves and Accuracy Ratios for Global Financial Institutions

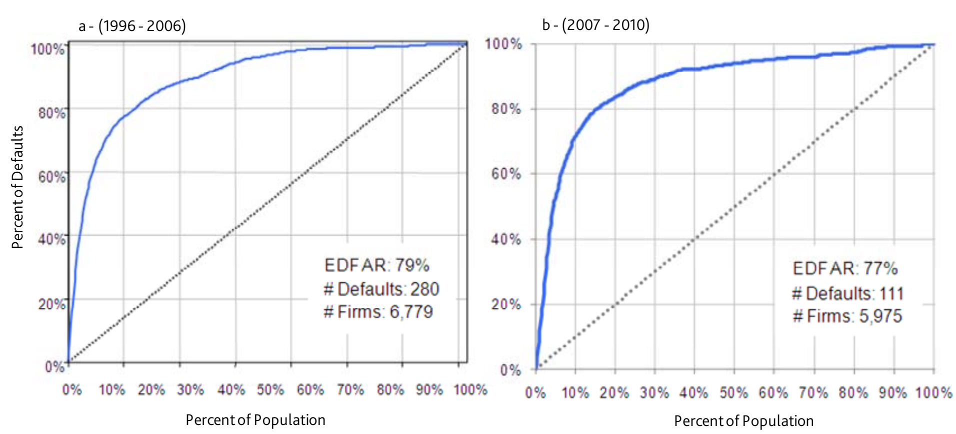

[4] Korablev and Qu (2009) and Gokbayrak and Chua (2009)

[5] CAP curves can be created for time horizons other than one year as well. The results in Figures 1a and 1b represent averages of cohorts formed at the beginning of each month for the historic periods cited.

[6] The operational definition of "default" is broad — bankruptcy is just one possibility. Entities are counted as defaulting if they are 1) formally declared in default or as bankrupt; 2) miss a scheduled interest or principal payment; and 3) engage in a restructuring/exchange of its securities that leaves creditors disadvantaged. Note, too, that default may occur for any credit instrument in a firm's capital structure. So the default of a bank's capital securities would be counted, even if senior creditors were guaranteed by the relevant government. We address this point at some length in the second half of the paper.

Accuracy Ratios are a concise way to capture the information represented by CAP curves. From Figures 1a and 1b it can be readily ascertained that EDF measures would do an even better job of signaling default risk if a greater percentage of defaulters had really high EDFs — for example, if 95% of the defaulting firms had EDFs in the top 10% of the sample, rather than 84%. This would push the CAP curve farther towards the upper left-hand corner of the plot. The Accuracy Ratio can be thought of as the surface area between the 45-degree line on the graph and the CAP curve, divided by the entire area above the 45-degree line.[7] But as many readers will realize, the Accuracy Ratio is really a correlation statistic, in that it measures the correlation between the ranking of the defaulters and that of the entire population.

The EDF Accuracy Ratios for the pre-crisis and crisis periods are 79% and 77%. While we cannot, as noted, exactly compare accuracy ratios across different portfolios, these statistics suggest that EDF credit measures performed at least as well during the financial crisis as in the period prior to it. They also show that the model's rank ordering power is quite good compared to those of other default risk measurement systems.[8] By comparison, Accuracy Ratios for debt ratings are usually in the low- to mid-70% range on average. (The offset for ratings is a reduced level of signal volatility, a topic for another day.)

In terms of rank ordering of default risk we can say confidently that EDF metrics on financial institutions performed well

So in conclusion, in terms of rank ordering of default risk we can say confidently that EDF metrics on financial institutions performed well, both before and during the credit crisis and recession. But success in rank ordering doesn't equate to success in other ways. Another question is whether the *levels* of the EDFs were correct. That is, did the passage of time reveal that the EDF measures of defaulted firms were too high (i.e., well in excess of their realized default rates) or too low (the opposite)? We address this question of *level validation* in the next section.

Figure 2a and 2b — Level Validation for Global Financial Institutions

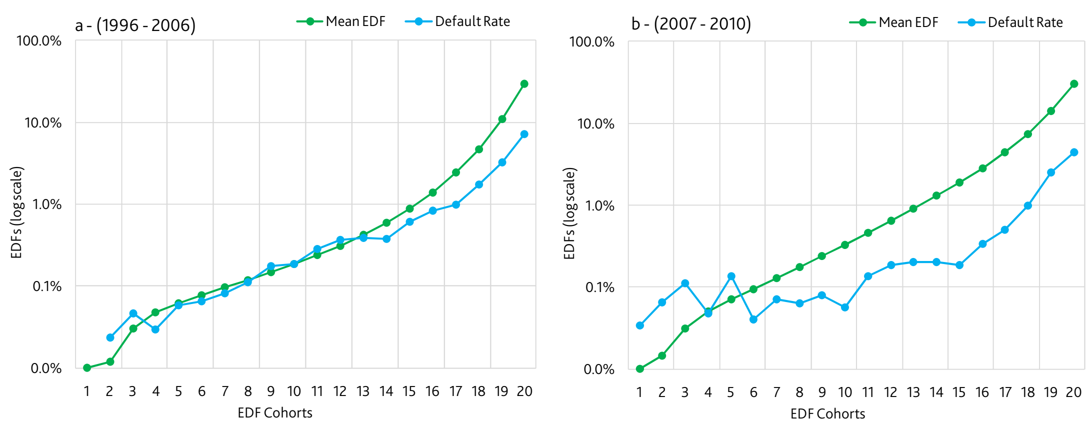

The evidence shows that EDF measures for financials weren't "too low" during the credit crisis and recession

Level validation — the challenge of getting it right

The second important performance measure of a credit risk management system is its level calibration: do predicted default rates correspond to actual, observed default rates? If the model is performing well, then realized default rates should be in line (in a statistical sense) with the EDF measures associated with them. Figures 2a and 2b compare the realized default rates for global financial institutions with their EDF measures on a one-year horizon basis, for the pre-crisis and crisis/recession period. In these graphs, firms are grouped into 20 equally sized buckets based on their EDF *levels*.[9] For each of the quantiles we calculate the realized default rates. For both periods the realized default rates are reasonably consistent with the levels of defaults predicted by the EDF measures. The anomalies on the left side of the graphs should not be surprising. Realized defaults among entities with low EDFs are quite rare, so variations of a couple of defaults can have big impacts on the realized rates. Thus, on aggregate, the evidence shows that EDF measures for financials weren't "too low" during the credit crisis and recession. To put this another way, the common impression

[7] The 45-degree line is important because if the CAP curve lay along it, then the model being tested would have no default ranking power. That is, for example, the top 20% of the sample by EDFs would contain 20% of the defaults.
[8] Bohn, Arora, and Korablev (2005), for example, compare pubic firm EDFs to PDs from a simple Merton model, credit ratings, and Z-score measures, and show that public firm EDFs exhibit higher rank ordering power over different samples and sample periods.
[9] As with Figures 1a and 1b, the results shown on Figures 2a and 2b represent averages of monthly cohort results.

that financial institution EDFs were low prior to the recession is correct. But the claim that they were too low relative to realized defaults is not supported by the evidence.

Indeed, the opposite is true, especially for the high EDF buckets. We see this on the right side of the Figures, where the EDFs are *above* the realized default rate. This is deliberate. Defaults are highly negative events for creditors, so a degree of conservatism is built into the model's calibration. Also, there is no central global repository of defaults — they must be compiled through a labor-intensive process. Despite our best efforts to record defaults — the EDF data base includes around 9,000 such events — we assume that we missed some. This means that the observed default rates (the blue lines) are almost certainly too low. Note, too, that the default rates in Figure 2b excludes some financials that were bailed out by governments. Arguably, these could have been included as defaulters, since they were only prevented from defaulting by extraordinary government actions. A final point concerns the fact that since the crisis began EDF measures for banks have been quite high. This has the effect of increasing the distance between the median realized default rate and the median EDF in Figure 2b, particularly for the high EDF buckets.

How to spot trouble early; focus on EDF performance vs. the peer group

While we can take comfort from the EDF level validation results shown in Figures 2a and 2b, it doesn't address a challenge faced by many users of EDFs; identifying the entities on which to focus. After all, many defaulted financials had low EDF metrics prior to their credit events — for example, the median EDF measure in the 15th bucket of Figure 2b is 2.8%. This is elevated, to be sure, but would not be a "red alert" for many users. So just looking for trouble among entities with high EDFs will not suffice.

The best risk screening strategy among low EDF entities is to compare the level and trend of a firm's EDF measure with that of its peer group

We've found the best risk screening strategy among low EDF entities is to compare the level and trend of a firm's EDF measure with that of its peer group. If the former is lagging the latter, then default risk could well be heightened. We see this in Figure 3, where the median EDF metric for firms that defaulted between 2007 and July 2010 (the red line) began to underperform the median for the entire data set as early as 2006. The sidebar on p.7 contains an example of this dynamic, for Lehman Brothers. But as noted at the outset, it's easy to find individual examples that support a point.

Figure 3 — EDFs for Financial Institution Defaulters vs. the Sector

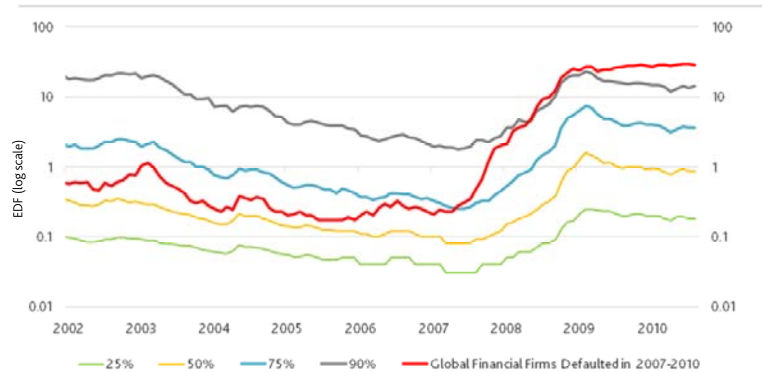

Figure 4 provides another illustration of the value of analyzing relative EDF performance. The horizontal axis measures the degree to which an entity's EDF *change* outperforms or underperforms the change of the median EDF for its sector over 12 months. The vertical axis shows the subsequent 1-year default rate for each relative change bucket. For example, entities whose EDF measures underperformed their sector medians by a factor of five (i.e., the firms' EDFs increased by five times the change of their sectors) suffered a subsequent average default rate of 5.2%. By contrast, those that underperformed by 2.5 times experienced a 2% average default rate.

Astute readers will realize that Figure 4 captures two concepts — the *movement* of an EDF vs. its sector and the *momentum* of an EDF — the idea that an event in one period (the relative rise of an EDF) is related to subsequent occurrences (an elevated default rate). We have found evidence of such momentum effects in previous studies of public firm EDF data and corporate bond and CDS prices.[10]

[10] Gokbayrak and Chua (2010) and Eckerstrom and Lam (2008)

Lehman Brothers; Lessons Learned

We can draw two lessons from Lehman Brothers' default. The first is the importance of focusing on the movement of a firm's EDF vs. its peers. Lehman's one-year EDF metric was 0.1% six months before the firm's default and 0.6% three months prior to the event. These figures are low in absolute terms, of course, but represented a significant underperformance vs. the metrics for global financial institutions, as the graph below illustrates. This is consistent with the users' tip on page 10 of tracking individual firm's EDF measures against those of their peers, and shows the value of EDF metrics in signaling deteriorating credit situations.

Secondly, Lehman is a useful reminder that EDFs reflect information obtained from equity markets and financial statements, and that they can be no better than the quality of this information. With the benefit of hindsight we know now that the markets, regulators, and others were operating with less than complete information about Lehman's true credit quality.

For example, the firm used certain accounting treatments to move liabilities off balance sheet on a temporary basis. As has been widely reported, Lehman sold equity and fixed income securities to a UK subsidiary and achieved "true sale" accounting treatment, even though they intended to repurchase these assets immediately after the close of their quarterly accounting dates. Inclusion of these "Repo 105 and 108" transactions into liabilities increased Lehman's leverage ratios, since about $60 billion was concealed that otherwise would have had to go onto the balance sheet. So from an EDF model perspective, LEH's liabilities and hence leverage were understated. Had these repos been reported correctly at June 30, 2008, the firm's EDF metric would have been higher.

Also, in the firm's last published financials (June 30, 2008) it reported an unencumbered liquidity pool of $35 billion - $40 billion. This was a big comfort factor for investors. The definition of unencumbered is that the assets are not pledged overnight. However, we subsequently learned that Lehman would use these assets to secure borrowings from their banks *intra-day*, but closed the transactions out each afternoon so that they were able to define the assets as unencumbered, (when for all practical purposes they were really pledged to the banks). Just a few days before their bankruptcy filing, the counterparty banks asked for more collateral margin from Lehman, and also asked for the pledges of the unencumbered assets to remain in force overnight. Thus, the "unencumbered asset "pool vanished. Word that the counterparties were asking for more collateral than Lehman could provide made its way out through the grapevine on the Street, hastening the company's downfall.

EDF metric for Lehman Brothers vs. its sector

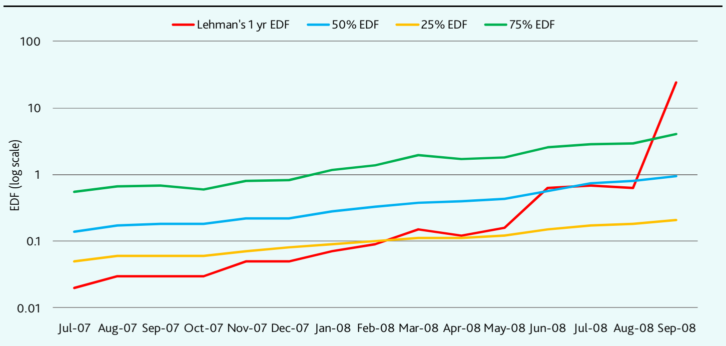

Figure 4 — Default Rates by Relative Performance (EDFs vs. Their Sectors) Bucket

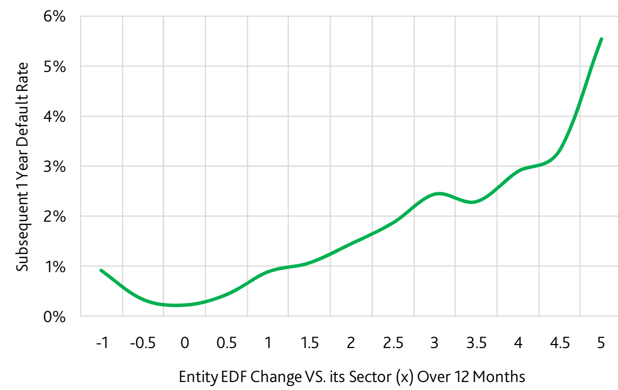

We now move from an analysis of historic EDF data to the present situation, specifically to the question of why many bank EDF levels are so high, and what users can (or should) do about it.

EDF Metrics and Banks: The Current Situation

Negative signals for banks

As we can see in Figure 5, since the credit crisis began A-rated financial institutions have exhibited higher median EDF metrics than comparably rated industrial companies. Moreover, a breakdown of median EDFs within the financial institutions space shows that for the past year banks have been the real culprits in this regard (Figure 6). As a result, the differential between banks' median Moody's rating and median EDF measure is quite large, when the latter is mapped to the Moody's rating scale (Figure 7).[11]

Figure 5 — Median EDF Metrics for Single A Entities by Sector

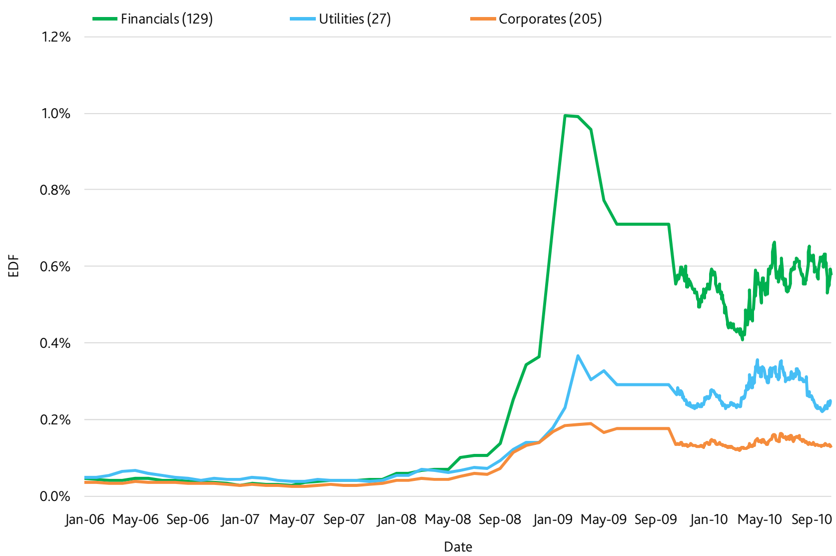

Figure 6 — Median EDFs for Financial Institutions Sub-groups

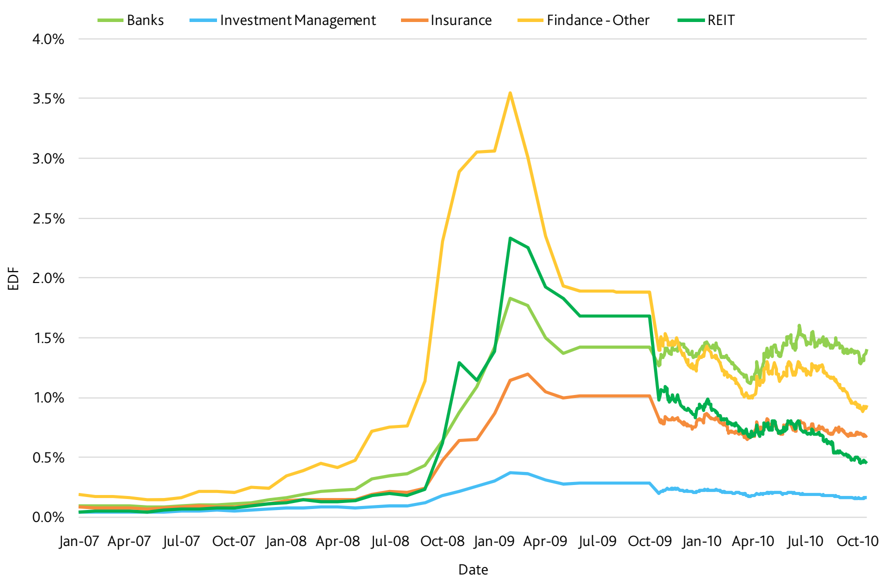

[11] The mapping from EDF measures to implied ratings is determined by median EDF measures of firms in rating classes using Moody's KMV "spot median" methodology. After calculating median EDF measures, the EDF range within a grade is computed from the median EDF of two adjacent rating grades. Once we have the EDF ranges for each category, we are able to assign an equity-implied rating for each EDF value.

Figure 7 — Financial Institutions Median EDF, EDF-Implied Rating and Moody's Rating (15-Day Moving Average)

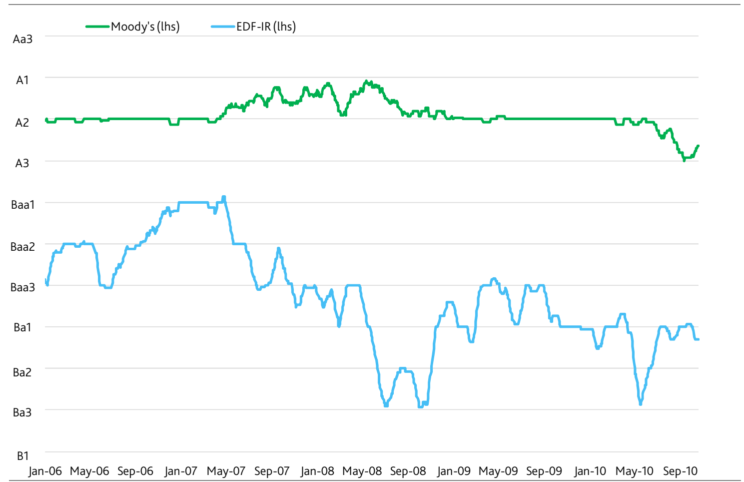

For many users, the elevated EDF levels for banks, as presented in Figures 6 and 7 are counterintuitive, which is a polite way of saying "hard to believe". We can think of examples like JP Morgan (Moody's senior rating of Aa3, EDF of 0.65%, EDF-implied rating of Ba3) or UBS (Moody's rating of Aa3, EDF of 2.22%, and EDF-implied rating of B3). The response that such EDF levels are "hard to believe" stems from users' long familiarity, and general comfort, with many large financial institutions. Also, ratings provide a broadly accepted framework for thinking about risk, so large deviations of banks' EDF metrics from their ratings often raise issues.

Different levels of the capital structure, different levels of risk

A bank's EDF measure could well overstate the risk to which investors and risk managers are typically exposed

The key driver of banks' seemingly high EDF measures is that *the model assesses default risk across a firm's entire capital structure*. This is of little significance for corporations, where most debt is at the senior level, but the implication for banks is profound. Bank liabilities range from senior obligations to "hybrid capital", such as preferred shares, that rank just above equity. As a result, a bank's EDF signal reflects the risk of its lowest-ranked obligation. The challenge is that for most investors and risk managers, the bulk of their exposures to banks are at the *senior* level. This is particularly true if we include counterparty and trading relationships, such as foreign exchange lines. As a result of all this, a bank's EDF measure could well overstate the risk to which investors and risk managers are typically exposed.

We note, too, that the impact of these capital structure considerations on EDF measures has become much bigger since the onset of the credit crisis. In many jurisdictions new laws and regulations have drawn a pretty clear line between bank senior obligations and those located further down the liability structure. The former are often eligible for some form of government support in the case of a crisis. The latter are not — or at best, the new rules are silent on the matter. In such cases, given the large number of unpleasant surprises over the past few years investors and risk managers are much inclined to assume the worst.

The greater gap between the haves and have nots (or to put it in a more precise form, obligations that will be supported, and those that probably won't, come crunch time), has also been recognized by the rating agencies. Figure 8 shows Moody's average bank senior and preferred share ratings. The differential is now seven rating notches, up from three notches early in 2007.

So far in this section we have confined our discussion to the varying probabilities of default associated with different levels of bank capital structures. This is only part of the story, as most readers will know. For many EDF users "risk" encompasses both default risk and loss-given default (combining these two concepts — PD*LGD — gives you the *expected loss rate*, of course). The point is that the lower level of risk on senior bank obligations (i.e., lower than signaled by EDF levels) can often be due to smaller levels of loss upon default. Government support of ultimate repayment on senior obligations (described further below) is an example of this. If a government takes over an insolvent bank it will strive to make senior creditors whole, but there could well be a delay in payment (which counts as a default) until things are sorted out.

There's an echo of this in ratings as well, in that Moody's ratings reflect expected credit loss. So if an obligation's expected LGD is very small, then it rating can be quite high even if its PD assessment is elevated.

Other obligations further down the capital structure might have the same PD assessment, but carry a much higher LGD estimate. As a result, their ratings would be lower. This is also a driver of the lower preferred share ratings that we see in Figure 8.

Figure 8 — Average Moody's Ratings for Bank Senior Debt and Preferred Stock

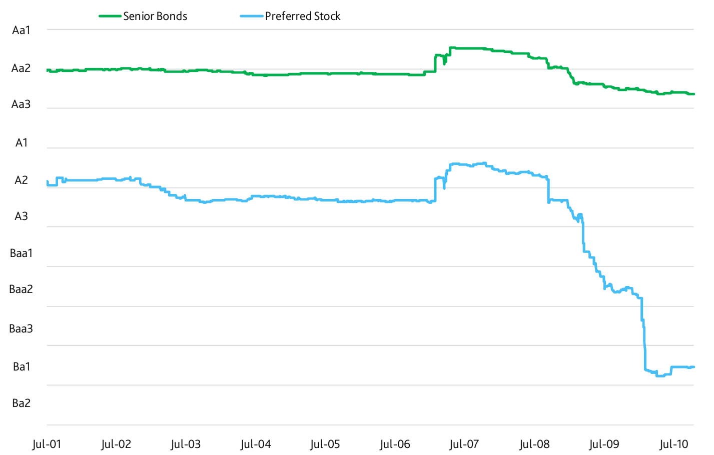

All this means that EDF measures on many banks "feel" elevated because they're focused on risk at the bottom of the capital structure, which is not where many counterparties have exposure. Moreover, the levels are often far removed from what would be suggested by the entities' senior ratings. This is at least partly explained by the fact that very low assumed loss-given default levels offset what might well be high default rate estimates.

A practical tip: focus on changes in EDF levels, and not just the EDF levels alone

We believe there is significant value in focusing on changes in EDF levels over a period of time

The question for EDF clients is what to do with this information. If their exposures differ, in level terms, from what is being captured by the EDF measures, how can they best make use of the product? In such cases we return to an earlier point: that there is significant value in focusing on changes in EDF levels over a period of time. An increase in a bank's EDF measure could well signal greater default risk, even if the current *absolute* level might be excessive compared to a user's actual exposure. In this way EDF metrics can provide early warning of impending difficulties. The approach is consistent with the tendency for changes in market-based risk signals to lead changes in fundamental credit assessments, such as ratings.[12] Figure 9 demonstrates this, with the average Moody's rating for banks being "pulled towards" the lower market levels over time.

Figure 9 — Banks: Moody's Sr. Rating, Bond-IR, CDS-IR and Equity-IR (15 Day Moving Average)

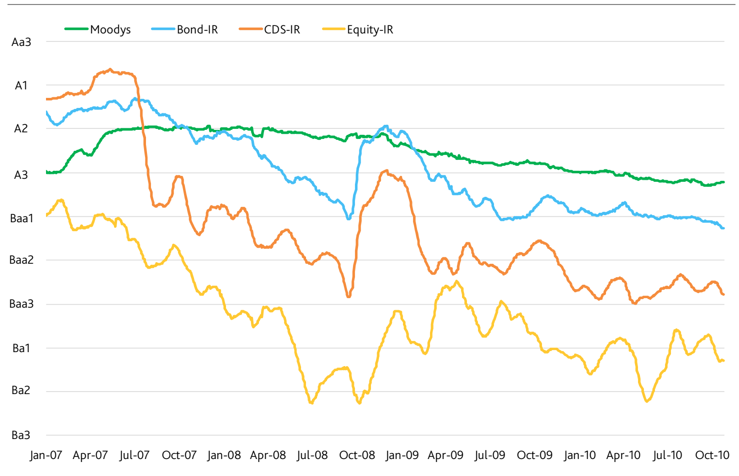

[12] Sun and Choi (2010).

We cannot exclude the possibility that the relatively high EDF metrics for many banks will prove, with the passage of time, to be correct

A second option is to use CDS-implied EDF measures as an alternative. This recently-launched EDF variant derives EDF measures from firms' CDS spreads.[13] Since the CDS contracts used in the model are written with reference to senior obligations, they eliminate the capital structure issues discussed above.

Should the EDF model be changed?

The foregoing discussion raises the question of whether the public firm EDF model should be changed so that it captures default risk at the top of a bank's capital structure, rather than at the bottom. The current incarnation is Model 8.0, so there have been many instances where insights gained into risk modeling, more extensive data, and evolving customer needs have led to model updates. This will not be one of them.

To begin with, we cannot exclude the possibility that the relatively high EDF metrics for many banks will prove, with the passage of time, to be correct. That is, the realized default rate will end up being in line with that signaled by the entities' elevated EDFs, within a given confidence interval.[14] After all, risks abound in the banking system. At this writing the European sovereign crisis is once again gathering pace; the US housing market shows no sign of a sustained rebound; and global economic growth is weak, with huge trade and payment imbalances among regions. Thus, the EDFs on banks shown in Figure 7 and elsewhere might well simply be reflecting an "inconvenient truth" — that the real outlook for the banking sector is worse than many believe. Indeed, as Figure 9 demonstrates, *all* market-based risk signals are more negative for banks than fundamentally-based analytical approaches such as Moody's ratings.[15]

Another point is that many investors own subordinated and hybrid debt. So confining the EDF signal to senior level risk would just create the opposite problem for them.

Finally, structural models of default, such as the EDF public firm model, typically do not allow for differentiation among risk levels in a firm's capital structure.

Banks' complex liability and corporate structures have further implications for investors and risk managers utilizing the EDF Public Firm model. We discuss these in the next section.

[13] For details, please see Dwyer *et al* (2010)

[14] This hypothesis has been confirmed statistically in Korablev and Qu (2009)

[15] The bond-, CDS-, and equity-implied ratings in Figure 9 are part of Moody's Market Implied Ratings platform, which provides spreads and EDF measures mapped to Moody's rating scale. For details please see Munves *et al* (2007).

Bank Structural Issues and Their Impact on EDF Measures

A question of structure

We start with holding company/operating company considerations.

EDF metrics reflect a blend of bank holding company and operating entity PD risk

A number of bank groups, particularly in the US, consist of holding companies and operating subsidiaries. The latter are the counterparties for many EDF users — meaning, for example, JP Morgan Bank NA, and not JP Morgan Chase & Co. The public firm EDF model calculates probabilities of default based upon consolidated financial statements, and the group's equity price and volatility reflects the performance of the consolidated entity. This means that EDF metrics reflect a blend of bank holding company and operating entity PD risk. The two cannot be separated by the model. One reason is that the group's share price reflects the performance of all group entities. And consistent with this, the model utilizes consolidated financial statements.

Bank holding companies are structurally subordinated to their regulated subsidiaries — the reverse of the usual arrangement for unregulated corporate groups. This is due to regulators' powers to "resolve" a potentially insolvent bank through a receivership or similar mechanism. Less drastically, regulators will limit payments from an operating company to a holding company, if they judge this necessary to preserve the former's financial strength. Either way, the aim of regulatory action is to protect small depositors, who are at the operating unit level. Such interventions often lead to the default of holding company obligations, as the issuer simply runs out of cash to pay interest and principal.

At this point many readers may be thinking "I get the point that the holding companies often default before their operating subsidiaries, but why does it matter? Won't the operating units be dragged into bankruptcy as well?" In fact, they usually aren't.

A key reason is that cross-default clauses are much less common for banks than for other types of non-sovereign issuers. These clauses, found in bond prospectuses, mean that the default of one debt issue triggers the default of all issues. However, in many jurisdictions, including the US, bank holding company and operating company [subordinated] debt issues do not have cross-default clauses. They also lack rights of acceleration for subordinated debt in the event of default. To top it all off, US bank bonds are exempt from registration requirements under the Securities Act of 1933. This means that they do not even provide prospectuses to potential investors — only offering circulars.

The end result is that bank operating units can and do continue as going concerns, even while their parents are in bankruptcy. For EDF users this provides another reason why, in some situations, EDF measures for banks can overstate their actual default risk.

We now move on to a consideration of bank deposits, and how EDF clients who hold these assets should think about the risk signals from the model. Again we are faced with a combination of both PD and LGD risks.

Not all liabilities are the same

A review of bank crisis resolutions around the world shows a clear pattern of regulators treating depositors differently than other creditors. Generally, in their striving to keep depositors from suffering losses regulators prefer to arrange for the assumption of deposits by a successor institution — even when some of the deposits are uninsured. National City Bank, Wachovia Bank, Colonial Bank, and Washington Mutual are recent US examples in which *all* depositors were made whole (however, other creditors at WAMU took losses). Outside the US, authorities in Hong Kong, Ireland, and New Zealand have expanded their deposit insurance schemes to offer more explicit protection against default.

Sometimes things don't work out so neatly — government seizure or receivership of a troubled bank has been known to result in a declaration of default on some or all non-guaranteed obligations, including deposits. So in such situations an elevated EDF metric would be correctly capturing default risk. However, government involvement means that the historic loss-given default experience at the bank deposit level has been *de minimus*, even for non-guaranteed obligations. In this scenario a depositors' ultimate loss risk is greatly reduced.

Moody's ratings reflect these varying risks of default and loss-given default; the difference between the deposit rating for JP Morgan Chase Bank (Aa1) and the preferred stock of JP Morgan Chase & Co. (A3) is five notches, for Bank of America the deposit rating (Aa1) is 10 notches above the holding company's (Bank of America Corp.) preferred stock rating (Ba2).

All in all, we can see why a rise in the bank's EDF measure as it becomes increasingly distressed might not reflect an increase in default risk, especially at the depositor level. And if default does occur, the ultimate loss could be quite low. But no creditor wants to be caught unawares, so the early warning properties of the EDF will still prove valuable. At the end of the day, what an EDF user does will depend on how he or she interprets the signal, and how they weigh it with other, more qualitative considerations.

In the next section we take a deeper dive into bank liability structures, in particular the impact on EDF metrics of reclassifying bank liabilities . We address the questions of bank capital securities (at the bottom of the liability structure), and deposits at the top.

Looking at the top and at the bottom

Bank capital instruments offer protection upon default to more senior creditors.[16] Indeed, they are specifically designed for this role, and are approved by bank regulators. Regulatory capital includes preferred shares and some types of subordinated debt, and can make up a significant share of a banking group's total capital. For example, JP Morgan Chase's Total Risk Adjusted Capital Ratio (including qualifying subordinated debt and preferred stock) is 15.5%, giving a leverage (or gearing) ratio of 6.5.[17] Taking common shareholders' equity alone gives an equity to asset ratio of 7.8% and leverage of 13 times. The Public Firm model counts only shareholders' equity as capital. This increases market leverage, raising bank EDF measures. We should also note that what classifies as equity reflects how banking company financial statements are reported — these are not modeling choices.

EDF users with exposures to senior bank obligations can recognize the quasi-equity benefits of bank capital securities by utilizing the solver feature of CreditEdge Plus

EDF users with exposures to senior bank obligations can recognize the quasi-equity benefits of bank capital securities by utilizing the solver feature of CreditEdge Plus™, the web-based platform through which most clients access EDF measures. In this exercise they would adjust a bank's liability structure to reclassify some or all of its capital securities as equity. Doing so would reduce the entity's market leverage and lower its EDF measure. We provide an example of such an exercise in Appendix 2.

We can also consider the case of deposits, which are the top of a bank's liability structure. Deposits are a significant and important source of funds for many banks — for example, JPMorgan Chase & Co.'s deposits of $903 billion constitute 43% of consolidated liabilities.

Deposits are unique in many respects, making them different than traditional corporate debt instruments. To begin with, as discussed a significant portion of bank deposits are insured by national regulators, making them an unusually stable source of funding, even in times of stress.[18] Related to this, banks can raise new deposits with implicit government backing to replace deposit outflows. While bank runs have caused many failures over the years, banks maintain a significant pool of primary liquidity and secondary liquidity to meet seasonal and unexpected outflows. Moreover, deposits often benefit from government-related back-up liquidity facilities that can provide cash to meet withdrawals. An additional feature of deposits is that unlike corporate debt, many types are primarily for transaction accounts. This means that they are non-interest bearing, and thus have no debt service cash flow needs.

We can combine these considerations with the previous points about deposit protection in bank liquidations to make an argument that deposits can be segregated from other liabilities. Clients can do so by using CreditEdge Plus' solver function. We perform such an exercise for JPMorgan in Appendix 3.

We emphasize that such adjustments are judgment calls on the part of EDF users. Deciding whether or not to make them is part of the qualitative credit analysis carried out by risk managers and investors. This is consistent with the point we made at the outset of the paper — that EDF measures are meant to improve clients' risk management practices, not to substitute for them.

Sensitivity of bank EDF measures to movements in stock prices

Finally, we address a different point — the special sensitivity of bank EDF measures to their equity prices.

[16] Some capital instruments also allow interest or dividend payments to be suspended without such actions being events of default.

[17] All data is as of September 30, 2010.

[18] The individual deposit guarantee limit varies by country. It is currently $250,000 in the US, for example.

All entities' EDFs react to movements in their stock prices.[19] In some cases, a lower stock price reduces the market value of an entity's assets, increasing its market leverage, and thus raising its EDF measure. This dynamic is especially powerful for banks, since their high market leverage places the market value of assets structurally close to the default point. The offset is that asset volatility, as calculated by the EDF model, is usually quite low. This in turn is due to the highly diversified nature of bank assets, and the fact that most are typically not marked to market.

Bank stock prices are depressed compared to pre-crisis levels. Some analysts might take the view that the equity markets have overshot to the downside, and that prices could rise in the next year. In such cases they can use CreditEdge plus' scenario calculator to recalculate a bank's EDF measure on this basis.

There are three channels through which an increase in a bank's stock price can cause its EDF level to change. These work in different directions.

First, a higher stock price increases the Market Value of Assets (this assumes that the other inputs are kept constant), decreasing the EDF. Second, a higher stock price raises the equity-to-asset ratio. Through the model's deleveraging calculation (see Appendix 1), this causes an increase in asset volatility and consequently increases the EDF. Thirdly, a higher stock price is typically accompanied by drop in equity volatility. This translates into lower asset volatility if everything else is held constant. The result is to offset the second effect, to some degree. That is, it lowers the EDF. The first effect typically dominates the second and third effects, meaning a bank's EDF metric level tends to decline as the value of equity rises. Appendix 3 works through these offsetting effects in another context.

In Figure 10 we take JP Morgan as an example of how changes in a bank's share price could impact its EDF metric. We simplify the exercise by holding the other EDF model inputs unchanged. But as noted above, this approach means that the EDF improvements due to increases in JPM's stock price are somewhat exaggerated.

The pre-crisis peak for JMS's common share price was $51.20 (on February 14, 2007). On that date the market price to book value ratio for JPM was 153% and the one-year EDF measure was 0.01%. At its low point (March 6, 2009) JPM stock closed at $15.03. Its market to book ratio was 44% and the one year EDF metric had worsened to 1.00%. Subsequently, the stock has recovered to a close of $38.15 on December 1, 2010, with a market to book ratio of 90% and a one year EDF metric of 0.64%. This EDF metric maps to an equity-implied rating of Ba3.

Figure 10 — The Relationship between JPM's Share Price and its one-year EDF (as of Dec. 1, 2010)

JPM Common Equity Price	One Year EDF Metric
$40.00	0.56%
$45.00	0.46%
$50.00	0.37%
$55.00	0.30%
$60.00	0.25%

[19] However, movements in EDF measures don't just reflect changes in stock prices. Illustrating this is the fact that the default risk identification power of the model is significantly greater than for equity returns alone (Sun, 2010).

Conclusion

Clients of the public firm EDF model are intensely focused on the metrics for banks, reflecting the importance of bank counterparty risk for most investors and risk managers, the stresses of the past few years, and the current elevated levels. In aggregate, EDFs for financial institutions were at appropriate levels prior to and during the credit crisis and subsequent recession. The current high readings reflect a range of considerations, including real risks around the sector and factors specific to banks. We suggest ways in which clients can interpret EDFs for banks, and make adjustments to EDF levels using the CreditEdge plus platform.

The authors would like to thank Jing Zhang, Douglas Dwyer, Zhao Sun, Lisa Hintz, and Rob Nagel for their valuable comments on earlier drafts of this paper.

Appendix 1: The EDF Public Firm Model

A company's stock price isn't an assessment of its creditworthiness

The equity market provides an estimate of the future prospects of each firm. The dividend discount stock valuation model is based on this premise, of course. However, a company's stock price isn't an assessment of its creditworthiness. We need only to think of a highly leveraged firm, which can have a great stock price performance while also carrying a significant risk of default. In this way stock prices differ from spreads in the credit default swap and corporate bond markets. Credit spreads contain factors other than default risk, for example, an assumed loss-given default rate and a general market price of risk. However, there can be no doubt that credit markets are highly focused on default, and that wider spreads equate to higher risk. By contrast, there is only a weak connection between company stock prices and realized default rates.[20]

The EDF Public Firm model builds on insights contained in the Black-Scholes-Merton (BSM) structural model of default risk, the original work on which dates from the 1970s.[21] Figure 11 provides a visual representation of the Public Firm model, and encapsulates the model's two main drivers.

Figure 11 - Default Process in the Structural Model

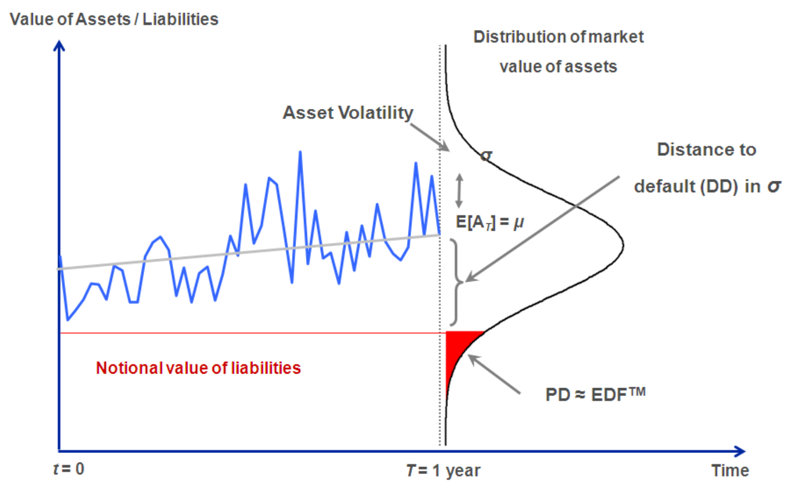

The model incorporates the components of classic credit analysis —business risk and financial risk — that are familiar to all practitioners

The first driver is the difference between the market value of the firm's assets and the book value of its liabilities at a future point in time.[22] In Figure 11, the future point in time equates to a one-year horizon (T=1); the expected mean level of the market value of assets is the grey line, and the book value of liabilities is the red line. The differential between the assets and liabilities is a measure of leverage, or financial risk, with a smaller differential equaling greater leverage, and thus more risk. The other driver is the expected asset volatility (represented by σ – the blue line is just one such possible path). Asset volatility is measure of business risk.

In this way the model incorporates the components of classic credit analysis — business risk and financial risk — that are familiar to all practitioners. In the Public Firm model, default is assumed to occur when the market value of a firm's assets falls below the book value of its liabilities, i.e., when it has negative net worth. The probability of this occurring is approximately equal to the shaded area in the right tail of the distribution in Figure 11. In the model the probability that the obligations will not be met is a function of the firm's distance to default (DD), which represents the number of standard deviations that the firm's asset value is away from the default point. DD can also be viewed as a volatility-adjusted market-based measure of leverage. DD is a crucial concept in the Public Firm model. Obviously, DD is reduced (and default risk

[20] Sun (2010)

[21] Merton (1974). See Crosbie and Bohn (2003) for a full description of the Public Firm model. Please also see Sellers and Arora (2004) for a description of model modifications made to account for financial institutions' special characteristics. Most of the model description in this Appendix first appeared in Dwyer & Qu (2007).

[22] The model holds the book value of liabilities constant over a one-year horizon. This assumption is relaxed when calculating EDF measures for longer terms.

increased) by either a rise in asset volatility or a reduction in the difference between the expected market value of assets and the book value of liabilities at the measurement horizon.

The EDF metrics calculated by the public firm model are calibrated to actual default rates. That is, unlike the BSM model, public firm EDFs reflect the empirical relationship between DD levels and historical default rates. This calibration, based on our empirical mapping, can have a large effect on estimated probabilities of default. For example, a firm with a DD of 4 has a near zero (0.003%, to be exact) likelihood of defaulting, whereas our empirical mapping indicates a 0.4% probability of default.

Structural models – a further discussion

Figure 11 and the accompanying text provides a useful conceptual framework for the public firm model. However, it doesn't touch upon one of the key aspects of structural models – the use on a firm's share price and share price volatility to calculate its market value of assets and asset volatility. We explore this in the following section.

Recall that a firm's share price and price volatility can be observed, while the market value of its assets and asset volatility cannot. Thus, the public firm model uses the Black-Scholes option pricing formula to derive a firm's *unobservable* firm values its *observable* equity values. The four inputs to this equation are a firm's debt payment, *D,* which we refer to as the **default point**, the **market value of the firm's assets**, *V*, **the volatility of assets**, σ_A, and the risk-free interest rate, *r*. The three bolded terms plus the DD are the model's key drivers.

A simplified example helps illustrate the linkage between a firm's unobserved asset volatility and market value of assets on the one hand, and its observed equity values on the other.

Let's take a company that has shares outstanding with a market value of *E*. Further, the company has a single class of debt, consisting of a zero coupon bond with a face value of *D*. As a zero coupon issue it has one payment (at maturity) at a fixed point in time, *T*. Owning the equity of this company is equivalent to holding a call option to buy the debt at a price of *D* (i.e., the face value of the obligation) with an expiration date of *T*. This reflects the fact that bondholders have a priority claim on a company ("priority" compared to shareholders, that is). So to fully own the company the shareholders must first pay off holders of the firm's debt. The value of the equity holders' call option will rise with the value of the firm. That is, the more the firm's assets are worth, the more advantageous it will be for the shareholders to pay off the debt and take full ownership of the company. As this Appendix describes, the Public Firm model transforms the value of the equity call option, and its volatility, into the market value of the firm's assets and asset volatility.

Owning a home with a mortgage provides a useful analogy. Even though it's common to say that you "own" your house, this isn't really the case, since you don't possess its title.[23] This is held by the bank that lent you the money. You only regain the title to the house (and thus truly own it) by paying off the mortgage. And the further the market value of the house is above the mortgage balance, the more likely you are to pay off the mortgage.

Option pricing theory tells us that a put and a call on an asset will have the same value if the strike price and exercise dates are the same, and volatility of the asset itself is the same. Thus, owning the firm's debt is equivalent to owning a risk-free bond with a face value of *D* at time *T,* and being short a put option on the stock in the company. The put has an exercise price of *D* and an expiration date of *T*.

This conceptual framework reflects that fact that if the value of the firm falls to zero, the shareholders will exercise their put by passing ownership of the firm to the bondholders. Note that they have this put option due to the limited liability nature of equity. The analogy of the homeowner is highly relevant here, given what's happened in the US residential real estate market over the past couple of years. In many jurisdictions homeowners with negative equity in their houses have the option to "walk away" from their mortgages, and the banks have no legal recourse to them. As has been widely reported, homeowners have exercised their put options by the millions.

Moody's Analytic's Public Firm model contains a number of important adaptations from the original Black-Scholes-Merton specification. The key developments are as follows

[23] Some readers may by unfamiliar with the term "title". This is the legal document proving ownership of a property such as a house or a car.

- The model incorporates several different types of liabilities, not just a single class of debt in the form of a zero coupon bond. Moreover, the model allows for default to occur at any point in time, and not just at a single point. This means it captures the failure by a firm to make interest payments.
- The model provides estimates of asset volatility for firms where only a limited time series of relevant equity prices are available. Such cases occur for newly public companies and for companies that undergo a merger, acquisition or spin-off that represents a substantial change in their business.
- The model calculates a consistent and credible measure of the default point by means of assumptions about debt maturities.
- A mapping of firms' DDs to observed default rates, which we have already noted. For firms with low DDs, actual default rates are much higher than would be implied from the cumulative normal distribution of asset returns assumed by many structural models. This is consistent with very bad asset returns occurring more frequently than would be implied by a normal distribution. The existence of "fat tails" in equity returns and bond returns is well documented in the finance literature, but less has been written on the distribution of asset returns.

In the following sections we take each of these considerations in turn.

Determining the value of a firm's assets and the default point

The standard Black-Scholes-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. One can then obtain an analytic expression for the value of equity by constructing a risk-free portfolio and solve the resulting partial differential equation using two boundary conditions relating to the value of the option at expiration and the value of the option when the underlying assets become worthless. In this context, the actual solution is the Black-Scholes option pricing formula.[24]

The structural model that forms the foundation for the public firm model is the Vasicek-Kealhofer (VK) model.[25] The VK framework assumes five types of claims on the firm's cash flow:

- Short-term liabilities
- Long-term liabilities
- Convertible securities
- Preferred stock
- Common stock

Incorporating these different types of contingent claims into the model changes both the asset value process and the boundary conditions.

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 the value of debt. For example, consider two firms with identical assets and debt but one pays a larger dividend. Obviously, 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.

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 holders of common stock. Consider two firms: A and B. Both have the same assets and debt. Company B, has a convertible security

[24] See Section III of Merton (1974). In the same paper, Merton develops several useful extensions to the "standard Merton model."

[25] Oldrich Vasicek developed the analytic model. Stephen Kealhofer led the research efforts that determined many of the empirical implementation details.

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 sold 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 amount of convertible securities, one will underestimate the market value of assets and overstate the probability of default.

Current and long-term liabilities

In structural models that use an absorbing default barrier, two approaches have been taken to defining 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 fall below a certain threshold, then the holders of equity chose to stop making payments on the debt and thereby turn over the firm's assets to debt holders.[26] In the VK model, the barrier is exogenous, in the sense that creditors put the debt back to the firm as soon the value of assets hits the value of the default point. The barrier is implemented as current liabilities plus a portion of long-term debt. We refer to this barrier as the default point. This barrier forms the second boundary condition of the model — the value of equity is equal to zero when the value of the firm's assets is equal to the default point.

Preferred stock

Preferred stock has both equity and debt characteristics, and the model takes both of these aspects into consideration. The VK model is able to incorporate various types of preferred stocks, including tradable preferred stocks and convertible preferred stocks. This completes the description of the different types of claims on the cash flows covered by the VK model.

Vasicek's model is an extension of other models that treat equity as a down-and-out perpetual option. Such extensions are necessary to credibly implement the structural modeling framework for actual firms. Actual firms have convertible securities, pay dividends, coupons and interest payments. These liabilities need to be explicitly modeled to generate a reasonable measure of default risk — particularly in the cases where these payments are unusually large.

In implementing the model, the default point is updated for every firm on a monthly basis based on publicly available information. For non-financial firms, the inputs are current liabilities and long-term liabilities less minority interest and deferred taxes.[27] There is often some ambiguity as to how to classify various items on a firm's balance sheet into long-term debt, long-term liabilities and current liabilities. Care is taken in making these distinctions. We refer back to original company reports and filings in addition to cross referencing the data against other sources when necessary. Given the default point, asset volatility and the risk-free interest rate, one can solve for the value of assets that sets the modeled value of equity equal to the observed value of equity.

Determining the volatility of assets

The Black-Scholes pricing formula was originally derived to value equity options with the volatility of equity being an important input. There is vast literature on different ways to model the volatility of equity. Extensions to the Brownian motion process include instantaneous jumps in equity value and time-varying stochastic volatility. One could implement these approaches directly from the observed time series of equity returns.

In structural models, one writes down an asset process and solves for the equity process as implied by the model. If the asset process has a constant volatility, volatility in the implied equity process is time varying. One implication of constant asset volatility is that as the value of assets falls close to the default barrier,

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

[27] For financial firms, these numbers are calculated from the balance sheet,. See Sellers and Arora (2004) for details.

equity volatility increases because of increased leverage. Consequently, de-levering a measure of equity volatility tends to understate the volatility for the firm's assets as asset value approaches the default barrier. Therefore, a direct measure of asset volatility is required, which is one of the key challenges in implementing a structural model: one needs to estimate the volatility of an unobserved variable. The VK model iteratively estimates both the value of a company's assets and its volatility.

Any estimate of asset volatility will be just that — an estimate. There are two sources of information available to estimate this volatility: information that is specific to the given firm (such as its equity price and liabilities history), and information for the population of comparable firms (their equity prices and liabilities history). We use a firm's specific information to estimate what we call empirical volatility, and we utilize the information from the population of comparable firms to estimate what we call modeled volatility. We combine the two. The weight on empirical volatility (relative to modeled volatility) is determined by the length of the time series of equity prices that is used in estimating empirical volatility.

The intellectual origins of combining empirical and modeled volatility can be found in Vasicek (1973). This paper develops a Bayesian approach to estimating β in the context of a Capital Asset Pricing Model. In this paper, rather than estimating a firm's β through simply the ordinary least squares (OLS) regression of the firm's returns on the market returns, the paper starts with a prior distribution of the firm's beta and then uses Bayesian methods to compute the posterior. The prior distribution of a firm's β is the distribution of β of comparable firms. Vasicek argues that the optimal estimate of β is the mean of the posterior distribution. Under this method, one combines the OLS estimate of β with the prior expectation of β to achieve an optimal estimation of β (in the Bayesian sense). When combining the two estimates, the weight on the OLS regression estimate increases with the number of observations used in the OLS estimate of β.

Empirical volatility

Given volatility and the default point, one can solve for the asset value that sets the value of equity implied by the VK model equal to the observed value. Thus for a given firm, one can construct a time-series of asset values. One can then compute what we call *empirical volatility* through an iterative method mentioned in *Crosbie and Bohn (2003)*: Using the VK model we compute a time series of asset values and hedge ratios from which we *de-lever* equity returns into asset returns. We compute the resulting volatility of asset returns, and then iterate until convergence. Remarkably, Duan, Gauthier, and Simonato (2004) showed that an iterative approach could be viewed as an application of the EM algorithm: an iterative procedure for estimating empirical volatility turns out to be the maximum likelihood estimate of asset volatility.

The iterative approach is superior to "solving for two equations and two unknowns" that was found in some academic and commercial implementations of structural models. As discussed in Ericsson and Reneby (2005), the issue with the two equations and two unknowns approach is that it assumes that equity volatility is a constant which is inconsistent with a structural model in which asset volatility is a constant. For companies with rapidly falling asset values, equity volatility is increasing rapidly due to changes in leverage. Consequently, asset volatility becomes understated for companies with rapidly increasing EDF credit measures. For the public firm model, the iterative approach grew out of poor results associated with attempts to de-lever equity volatility in the early 1990s.

A large corporate transaction — such as a merger or acquisition — may result in a permanent change in the volatility of the company. Post-event asset returns are likely to be more informative than pre-event asset returns for estimating empirical volatility. Consequently, empirical volatility is computed differently in the model following a corporate event. If there has been a large change in firm size or capital structure, then empirical volatility is computed by placing a larger weight on the post-event asset returns.

Modeled volatility

When the equity price history is limited we rely more on modeled volatility. Using empirical volatility for public firms as the dependent variable, we estimate volatility on the basis of the size, industry and geography of the firms, as well as some accounting ratios. For a new firm, the volatility used in the model is heavily based on modeled volatility. For firms with a long history of equity prices, the volatility used in the model is a combination of modeled volatility and empirical volatility.

In industry and geography aggregates, one observes empirical volatility changing over time in response to changing business conditions. Each month, modeled volatility is recalibrated so that on average modeled volatility is equal to empirical volatility. In this way, modeled volatility neither increases nor decreases changes in aggregate volatility that may occur as the result of changing business conditions.

Non-Gaussian relationship between Distance-to-Default and the EDF value

As the VK model is a barrier model, the model relates the asset value, the default point and volatility to the default probability via a first passage through time formula. Vasicek has noted that the probability of default for a first passage through time model is approximately equal to:

$$2 \times \Phi(-DD) \tag{1}$$

where DD is the so-called Distance-to-Default and Φ is the cumulative normal distribution. Distance-to-Default can be defined as:

$$DD\left(V, X_T, \sigma_A, T, \mu, a\right) = \frac{\log\left(V / \left(X_T + aT\right)\right) + \left(\mu - \frac{1}{2}\sigma_A^2\right)T}{\sigma_A \sqrt{T}} \tag{2}$$

Where *V* is the value of a firm's assets, X_T is the default point to the horizon, μ is the drift term, σ_A is the volatility of assets, *T* is the horizon and *a* represents cash leakages per unit time due to interest payments, coupons and dividends. The value of the firm's assets and volatility is computed as described above. The default point is computed as current liabilities plus a portion of long-term debt. For longer horizons, a larger portion of long-term debt is included in the default point to reflect that long-term debt becomes more important at longer horizons. Note that the DD varies considerably with the horizon under consideration. At longer horizons, the weight on volatility increases relative to the default point.

Empirically, there is a strong relationship between DD and the observed default rates — firms with a larger DD are less likely to default. Nevertheless, the actual default rate found in the data differs from the literal predictions of the model. Taken literally, the Brownian motion assumption on asset value implies a Gaussian relationship between DD and the EDF credit measure. Specifically, for a DD greater than 4, a Gaussian relationship predicts that defaults will occur 6 in 100,000 times. This would lead to one half of actual firms being essentially risk-free. This implication is not found in the data. Consequently, when implementing the model, we depart from the Gaussian assumption by implementing an empirical mapping.

Appendix 2: JP Morgan Chase: Adjusting Its EDF Metric for the Impact of Regulatory Capital and Deposits

Banks have complex liability structures, with different levels having different degrees of risk. We can take the ratings on JP Morgan Chase & Co. (JPM) as an example: its commercial paper rating is P-1, its senior unsecured rating is Aa3, its subordinated debt rating is A1, and it has a Baa1 preferred stock rating.

While there is differentiation within the EDF public firm model for short term and long term debt, there is no way to discriminate based on the seniority of the debt and other risk factors. However, these differences are recognized — indeed, they are created — by bank regulators. Among other distinctions, regulators count non-cumulative preferred shares and hybrid securities as Tier I Regulatory Capital (within a formula), and cumulative preferred shares and subordinated debt as Tier 2 Regulatory Capital.

Adjusting the EDF metric for JPM to reflect regulatory capital

As we noted in the body of the paper, many EDF clients have exposures to senior bank liabilities. In such cases, they may wish to give credit for the loss-absorbing nature of some bank capital securities. They can do so by using the Solver function in the CreditEdge platform to reclassify some or all of the bank capital liabilities as equity.

Figure 12

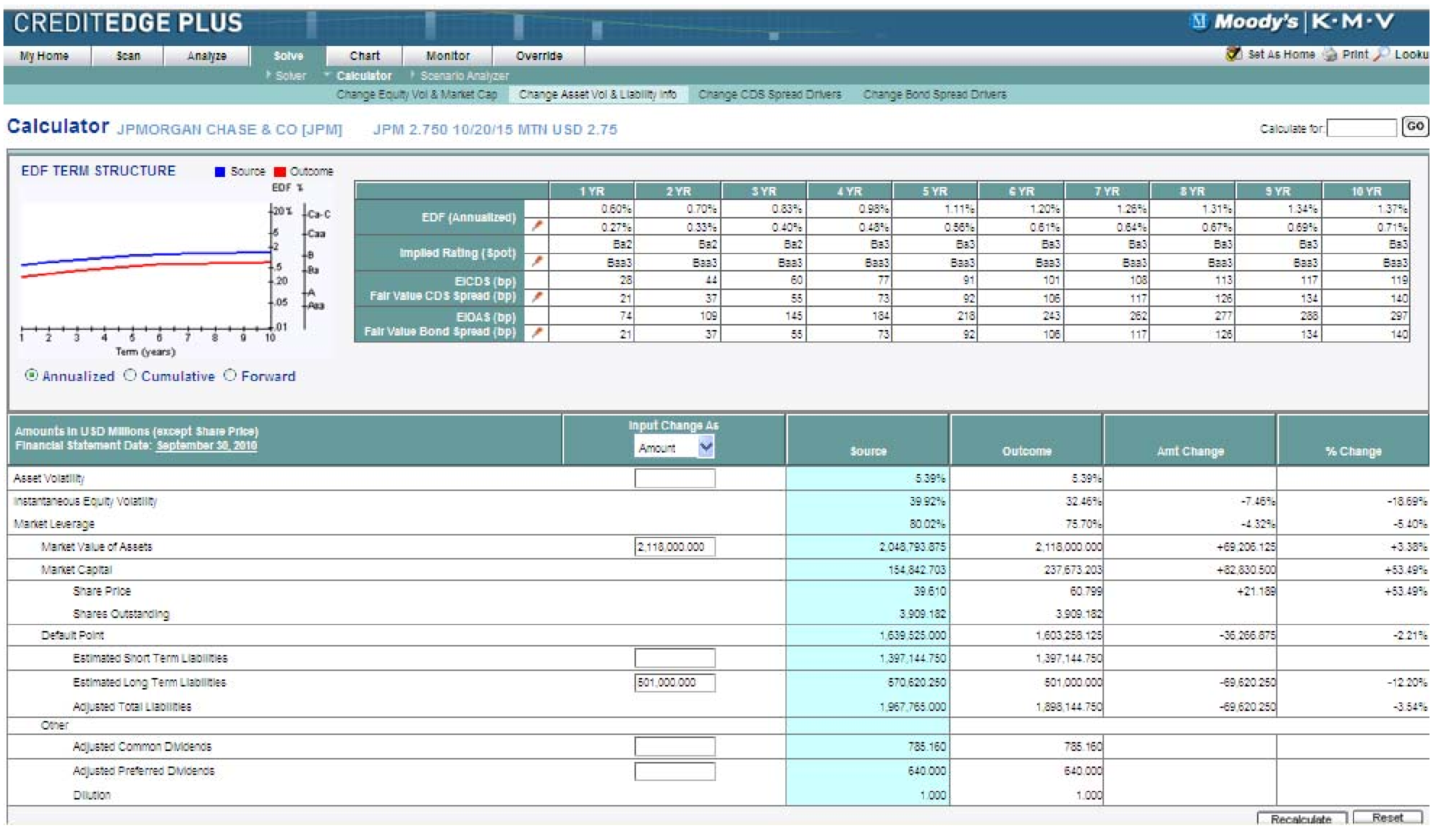

		1 YR	2 YR	3 YR	4 YR	5 YR	6 YR	7 YR	8 YR	9 YR	10 YR
EDF (Annualized)		0.60%	0.70%	0.83%	0.98%	1.11%	1.20%	1.26%	1.31%	1.34%	1.37%
		0.27%	0.33%	0.40%	0.48%	0.56%	0.61%	0.64%	0.67%	0.69%	0.71%
Implied Rating (Spot)		Ba2	Ba2	Ba2	Ba3	Ba3	Ba3	Ba3	Ba3	Ba3	Ba3
		Baa3	Baa3	Baa3	Baa3	Baa3	Baa3	Baa3	Baa3	Baa3	Baa3
EICDS (bp)		28	44	60	77	91	101	108	113	117	119
Fair Value CDS Spread (bp)		21	37	55	73	92	106	117	126	134	140
EIDAS (bp)		74	109	145	184	218	243	262	277	288	297
Fair Value Bond Spread (bp)		21	37	55	73	92	106	117	126	134	140

Amounts in USD Millions (except Share Price) Financial Statement Date: September 30, 2010	Input Change As Amount	Source	Outcome	Amt Change	% Change
Asset Volatility		5.39%	5.39%		
Instantaneous Equity Volatility		39.92%	32.46%	-7.46%	-18.69%
Market Leverage		80.02%	75.70%	-4.32%	-5.40%
Market Value of Assets	2,118,000.000	2,048,793.875	2,118,000.000	+69,206.125	+3.38%
Market Capital		154,842.703	237,673.203	+82,830.500	+53.49%
Share Price		39.610	60.799	+21.189	+53.49%
Shares Outstanding		3,909.182	3,909.182		
Default Point		1,639,525.000	1,603,258.125	-36,266.875	-2.21%
Estimated Short Term Liabilities		1,397,144.750	1,397,144.750		
Estimated Long Term Liabilities	501,000.000	570,620.250	501,000.000	-69,620.250	-12.20%
Adjusted Total Liabilities		1,967,765.000	1,898,144.750	-69,620.250	-3.54%
Other					
Adjusted Common Dividends		785.160	785.160		
Adjusted Preferred Dividends		640.000	640.000		
Dilution		1.000	1.000		

We can use JPMorgan to illustrate this principle. As of September 30, 2010, JPM had as a part total of Tier 2 capital $41 billion of preferred securities, and preferred securities and hybrids counting as Tier 1 capital of $28 billion, for a total of $69 billion.[28] Using the Solver function, we deducted $69 billion from long term liabilities and adding $69 billion to market capitalization. All other model inputs, such as asset volatility, were left unchanged. As shown in Figure 12, with these adjustments the one year EDF metric for JPM falls from 0.60% to 0.27%. This increases the corresponding EDF-implied rating from Ba2 to Baa3, only two notches below JPM's Baa1 preferred stock rating. Investors can use the Solver function within the model to recalculate an EDF metric depending on whether they wish to adjust liabilities and capital for some or all of the instruments that count as regulatory capital.

[28] In this exercise, we did not add back goodwill of $46 .7 billion nor AOCI (accumulated other comprehensive income) of a negative $2.9 billion which were both deducted from Tier I capital in JPM's calculation. Source: JPM 10Q

Appendix 3: Treatment of Bank Deposits in the EDF Model

As noted in Section 4, deposits are different structurally and legally from other types of liabilities. [29] However, because EDF metrics are based on consolidated financial statements, deposits are treated the same way as senior non-depository debt, subordinated debt and preferred stock.[30] In this Appendix we discuss how alternative treatment of deposits can result in lower EDF metrics for banking companies for users who want to make such adjustments.

We utilize the EDF Solver scenario calculator in the CreditEdge plus platform to assess the impact of eliminating deposits from JPM's consolidated balance sheet: One effect of this adjustment is to isolate default risk at the stand alone holding company. Consolidated deposits totaled $903 billion on September 30, 2010. We realize that eliminating all deposits is an extreme case, and does not discriminate between deposit classes (e.g., time versus demand deposits, insured versus uninsured deposits, or purchased deposits versus core deposits). Obviously in a real world sense, a deposit-taking institution would not pay off all its deposits, even if it had sufficient short term liquid assets to do so. But in taking this approach our intent is to demonstrate the relative importance of the deposits in calculating a bank's EDF measure. Each user will make his or her own decision on how to treat deposits.

The reduction of JPM's short term liabilities by $900 million (shown in Figure 13) was offset by a reduction in the market value of assets by the same amount. In this first example asset volatility[31] (which is usually lower for financial institutions than for corporate entities) was held unchanged at 5.39%. The reductions in assets and liabilities lowered the default point for JPM from $1.63 billion to $ 775 million, which resulted in a decline in market leverage, i.e., financial risk. Following our adjustments, JPM's one-year EDF metric fell from 0.60% to 0.10%, improving in its EDF-implied rating from Ba2 to A1.

As many EDF users know, asset volatility would be likely to rise if a substantial portion of liabilities were to be paid off (the model mechanics behind this are discussed in the next section). Thus, using an unchanged asset volatility of 5.39% is a "best case outcome". This means that a revised EDF metric of 0.10% represents the bottom of the possible range of EDF measures from the depository adjustment, since asset volatility would not fall below the 5.39% reported in the original unadjusted case.

[29] All bank holding company liabilities are structurally subordinated to bank level liabilities. Moreover, national depositor preference statutes in the US and in many other venues give all bank level depositors (insured and uninsured) preference in payment upon liquidation to all other senior creditors. We also note that the FDIC issued a Notice of Proposed Rulemaking for comment in November 2010, which would allow them to discriminate between creditors of the same legal seniority in the event of a receivership or bankruptcy (e.g., short-term versus long-term creditors within the same class of obligations might receive different treatment depending on the FDIC's final ruling.)

[30] The different expected losses for these liabilities is demonstrated by the Moody's credit ratings which rank preferred stock five notches lower than deposits for JP Morgan and 10 notches lower for Bank of America.

[31] In a general sense, asset volatility is a measure of the business risk of the firm. Technically, it is the standard deviation of the annual change in the market value of the assets, and is expressed in percentage terms of asset value. The higher the asset volatility, the less certain investors are about the market value of the firm, and the more likely the firm's value will fall below its default point. Please see Appendix 1 for details.

Figure 13

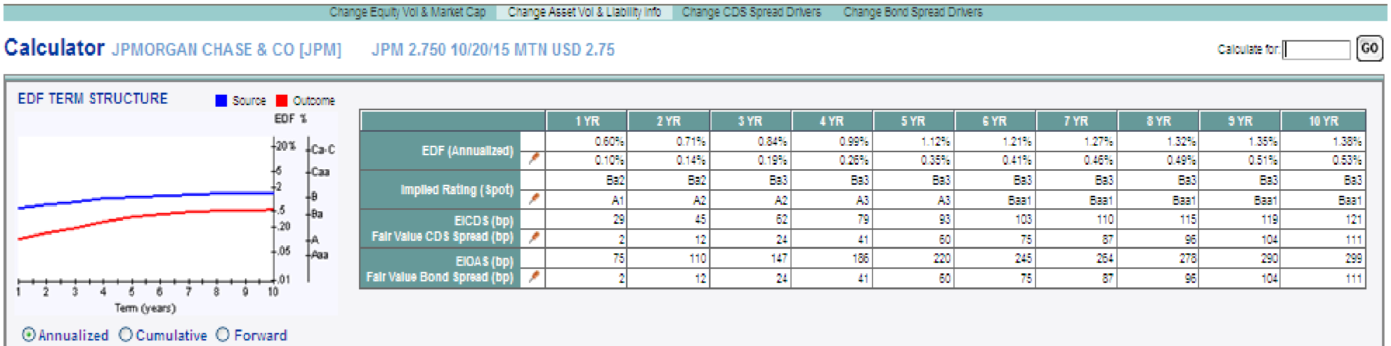

		1 YR	2 YR	3 YR	4 YR	5 YR	6 YR	7 YR	8 YR	9 YR	10 YR
EDF (Annualized)		0.60%	0.71%	0.84%	0.99%	1.12%	1.21%	1.27%	1.32%	1.35%	1.38%
		0.10%	0.14%	0.19%	0.26%	0.35%	0.41%	0.46%	0.49%	0.51%	0.53%
Implied Rating (Spot)		Ba2	Ba2	Ba3	Ba3	Ba3	Ba3	Ba3	Ba3	Ba3	Ba3
		A1	A2	A2	A3	A3	Baa1	Baa1	Baa1	Baa1	Baa1
EICDS (bp)		29	45	62	79	93	103	110	115	119	121
Fair Value CDS Spread (bp)		2	12	24	41	60	75	87	96	104	111
EIOAS (bp)		75	110	147	186	220	245	264	278	290	299
Fair Value Bond Spread (bp)		2	12	24	41	60	75	87	96	104	111

Amounts in USD Millions (except Share Price) Financial Statement Date: September 30, 2010	Input Change As Amount	Source	Outcome	Amt Change	% Change
Asset Volatility		5.39%	5.39%		
Instantaneous Equity Volatility		40.03%	32.21%	-7.82%	-19.54%
Market Leverage		80.08%	67.53%	-12.55%	-15.67%
Market Value of Assets	1,147,396.625	2,047,396.625	1,147,396.625	-900,000.000	-43.96%
Market Capital		154,060.797	92,059.180	-62,001.617	-40.24%
Share Price		39.410	23.549	-15.861	-40.25%
Shares Outstanding		3,909.180	3,909.180		
Default Point		1,639,525.000	774,802.500	-864,722.500	-52.74%
Estimated Short Term Liabilities	497,144.750	1,397,144.750	497,144.750	-900,000.000	-64.42%
Estimated Long Term Liabilities		570,620.250	570,620.250		
Adjusted Total Liabilities		1,967,765.000	1,067,765.000	-900,000.000	-45.74%
Other					
Adjusted Common Dividends		785.160	785.160		
Adjusted Preferred Dividends		640.000	640.000		
Dilution		1.000	1.000		

EDF users might consider selecting a higher asset volatility, which would reduce the impact of removal of deposits on the EDF metric. For example, they can use an asset volatility of 6.08%, which is derived from the average instantaneous equity volatility of 38 large comparable banking franchises with similar market leverage to JPM. In this case the one-year EDF metric would improve to 0.14% (from the original 0.60%) which would correspond to an EDF-implied rating of A3.

Offsetting effects

As we discuss in Section 4 of the paper, changes in any single EDF model input will likely result in changes to other inputs as well. When the direction of the EDF metric changes brought about by these inputs differ, the change in an EDF metric will be dampened by the offsetting directional changes of other factors. For example, a reduction in the EDF metric from removing deposits from a bank's consolidated liabilities and its market value of assets could be mitigated by changes to other model inputs such as asset volatility. It is reasonable to assume that JPM's business risk (expressed as asset volatility) could shift if deposits were eliminated. Users can make their own determinations of the appropriate adjustments to model inputs.

We have presented above the example of removing deposits from JPM's bank holding company's consolidated balance sheet with no other changes to the MKMV model inputs. The direct impact of removing deposits from the bank's liabilities is to lower its default point, causing a decrease in the bank's EDF metric. However, while this shrinks the size of JPM's balance sheet, it increases the consolidated *equity to asset ratio*. Within the EDF model asset volatility and the equity-to-asset ratio are related by the following formula $\sigma_A = \frac{\sigma_E}{\Delta E / \Delta A} \bullet \frac{E}{A}$, where E/A is the equity-to-asset ratio, (σ_E) and (σ_A) are equity and asset volatilities respectively, and $(\Delta E / \Delta A)$ is the "hedge ratio", measuring the sensitivity of equity change to asset value change.

An increase in the equity-to-asset ratio will cause an increase in asset volatility, which in turn will increase JPM's EDF metric

Assuming the first factor of the product $\left(\frac{\sigma_E}{\Delta E/\Delta A}\right)$ remains constant, an increase in the equity-to-asset ratio will cause a rise in asset volatility, which in turn will increase JPM's EDF metric. So removing the deposits will induce two offsetting effects: the lower default point will decrease the EDF measure, but the higher asset volatility will increase it — or at least minimize its decline.

Typically, the impact of a reduced default point dominates that of lower asset volatility, so the net effect is a drop in the banking company's EDF measure — but not as much as the case where asset volatility is artificially kept constant.

In the case of JPM, when its deposits are removed from its balance sheet as model adjustments, the equity-to-asset ratio increases from 19.88% to 32.41%. Without adjusting its impact on asset volatility JPM's EDF metric would drop from 0.60% to 0.10%. But with asset volatility adjustment, (which we have computed to be 8.79% using the above formula) JPM's EDF measure would only drop to 0.50%. We can consider 8.79% as the upper end of the range of possible asset volatilities in this scenario within the Public Firm model.

In addition, the user may also want to analyze the impact of capital structure change on the hedge ratio in order to fully assess its impact on the EDF metric. When the equity is viewed as a call option on the underlying firm assets (see Appendix 1), increasing the equity-to-asset ratio by removing the deposits tends to move the call option deeper in the money, causing an increasing in the hedge ratio, since option value is a convex function of the underlying asset. Given that the hedge ratio works as the denominator of the product, this in turn counter-reacts to the direct impact of equity-to-asset ratio on asset volatility. So the "true" asset volatility should be lower than what it would be without accounting for the impact on hedge ratio. This means a fully adjusted EDF level would be less than 0.50%, the value provided above. Thus, our analysis would calculate the one year EDF metric for JPM after its deposits are removed as ranging between 0.10 % and 0.50 %, (or an EDF-implied rating between A3 and Ba2), depending on the asset volatility an EDF user deems appropriate.

A final analytical point is that when the counterparty (or security issuer) is a bank holding company, investors should recall that the deposits reside within bank operating subsidiaries, and that *consolidated* financial statements can present a different picture of the liabilities versus that of the stand-alone holding company.

References

Black, Fischer and John C. Cox, "Valuing Corporate Securities: Some Effects of Bond Indenture Provisions", Journal of Finance, Vol. 31, No. 2, (May, 1976)

Bohn, Jeff, Arora Navneet and Korablev Irina, "Power and Level Validation of the EDF Credit Measure in the US Market" (March, 2005)

Crosbie, Peter and Bohn Jeff, "Modeling Default Risk", (December, 2003)

Duan, J., Gauthier G., and Simonato J., "On the Equivalence of the KMV and Maximum Likelihood Methods for Structural Credit Risk Models," (2004)

Dwyer, Douglas and Kocagil Ahmet, "Moody's KMV RiskCalc™ V3.1 United States", (June, 2004)

Dwyer, Douglas and Shisheng Qu, "EDF 8.0 Model Enhancements", (January, 2007)

Dwyer, Douglas, Zan Li, Shisheng Qu, Russell Heather and Zhang Jing, "CDS-implied EDF™ Credit Measures and Fair-value Spreads", (March, 2010).

Eckerstrom, Robert and Lam Chris, "Differentiating Risk Through Momentum Signals", (February, 2009)

Ericsson, Jan and Reneby Joel, "Can Structural Models Price Default Risk?", (January, 2005)

Gokbayrak, Ozge and Lee Chua, "Validating the Public EDF™ Model During the Credit Crisis in Asia and Europe", (November, 2009)

Gokbayrak, Ozge and Lee Chua, "Using EDF Momentum to Provide Insights into Credit Migration and Default Risk", (July, 2010)

Korablev, Irina and Qu, Shisheng, "Validating the Public EDF™ Model Performance during the Credit Crisis", (June, 2009)

Leland, Hayne, "Optimal Capital Structure", Journal of Finance, (September, 1994)

Merton, Robert C., "On the Pricing of Corporate Debt: The Risk Structure of Interest Rates", Journal of Finance, Vol. 29, No. 2, (May, 1974)

Munves, David W. , CFA, Hamilton T. David, Ph.D, "Moody's Market-Implied Ratings Description, Methodology, and Analytical Applications", (December, 2007)

Sellers, Martha and Arora Navneet, "Financial EDF Measures A New Model of Dual Business Lines", (August, 2004)

Vasicek, Oldrich, " A note on using cross-sectional information in Bayesian estimation of security betas", Journal of Finance, Vol. 28, (1973)

Sun, Zhao, PhD, CFA, "An Empirical Examination of the Power of Equity Returns vs. EDFs™ for Corporate Default Prediction", (January, 2010)

Sun, Zhao, PhD, CFA, Choi Yukyung and Hamilton T. David, Ph.D, "Market-Implied Ratings Transitions: Performance during the Credit Crisis", (June, 2010)

This page intentionally left blank.

To order reprints of this report (100 copies minimum), please call 212.553.1658.

Report Number: 129462

Author
David W. Munves, CFA 1.212.553.2844
david.munves@moodys.com

Contact Us
Americas : 1.212.553.4399
Europe: +44 (0) 20.7772.5588
Asia: 813.5408.4131