July 22, 2010

Ruben Hovhannisyan
Vice President
U.S. Fixed Income

Understanding Fixed Income

Quantitative Credit Analysis: An Introduction to Mertonian Models

Introduction

Corporate debt is one of the major sources of financing used by companies around the globe. U.S. corporate bonds represent one of the largest asset classes in the fixed income universe with approximate value of $6.9 trillion as of December 31, 2009, according to the Securities Industry and Financial Markets Association. The colossal size of the corporate bond market underlines the significance of fully understanding the properties and limitations of the models used for managing credit risk and evaluating corporate debt. Hence, it comes as no surprise that credit risk modeling has been one of the central topics in financial research and literature for more than four decades now. Despite the tremendous progress made in the field, the Great Recession, which brought the world's financial system to its knees in 2008, highlights the imperfections inherent in some financial models designed to manage corporate debt risk. Indeed, it is interesting to note that the event most emblematic of the crisis, the failure of Lehman Brothers, was barely on the radar screens of these quantitative models as recently as two weeks prior to the filing. That said, quantitative modeling of credit risk has a large following and notwithstanding some significant limitations, can provide valuable insight to the credit investor or risk manager.

This review offers a brief introduction to credit default models and provides the basic insights behind the widely-used framework to evaluate corporate debt default risk introduced by Robert Merton in the early 1970s.

Evolution of Credit Default Models

The first attempt to predict and quantify credit default probability dates back to the 1966 publication of William Beaver's univariate financial distress model. Beaver applied statistical techniques to evaluate the ability of various single accounting ratios to predict bankruptcy of a firm. Recognizing that the financial health of a company is better assessed by evaluating multiple key ratios simultaneously, Edward Altman further applied a discriminant analysis method to combine five financial ratios (using variables from the income statement and balance sheet) to produce a single number, called the Z-Score, which would predict the probability of a given company going bankrupt.

In 1974, less than 10 years after Altman's Z-score model was developed, Robert Merton's seminal work on corporate debt valuation introduced a completely new approach to pricing corporate bonds and evaluating credit default risk. Merton's model, variously called a structural credit model or contingent claims model, postulates that owning a corporate bond is analogous to simultaneously holding a long position in risk-free debt and a short put option on the assets of the firm issuing the debt. The critical insight of this approach is the recognition that a corporation's equity holders possess the ability to file bankruptcy and, in so doing, effectively "put" the company back to the bondholders. As such, being long a corporate bond inherently means that the bondholder has shorted a put on the company assets to the stockholders. The model further incorporates the Black-Scholes option pricing framework to calculate the value of corporate debt. Additionally, Merton combined firm asset value, debt value, and asset volatility to derive a default probability measure called "distance-to-default."

Although Merton's model did not take the world by storm immediately after being published, the model and its various extensions, aimed at relaxing some of the simplifying assumptions inherent in the original model, later became popular tools used for analyzing credit default probabilities.

Merton's Insightful Observation

Mertonian models are predicated on the basic idea that a firm's assets are jointly owned by equity and bond holders of the firm, where bond holders have a fixed claim on the firm's assets while equity holders' claim is contingent upon the value of the firm's assets relative to the value of debt. If the value of the assets is greater than the value of debt, then equity holders have a claim on the difference. Conversely, if the value of the assets is not sufficient to pay down the debt at maturity, the firm will default on its debt and the bondholders will take the ownership of the firm's assets and liquidate them in order to recover at least part of the defaulted payment. Recognizing this basic dynamic inherent in the capital structure of a firm led to Merton's insight that being a stockholder is analogous to being long a call option on the assets of the firm with a strike price equal to the face value of the debt.

Equity Holders' Contingent Claim

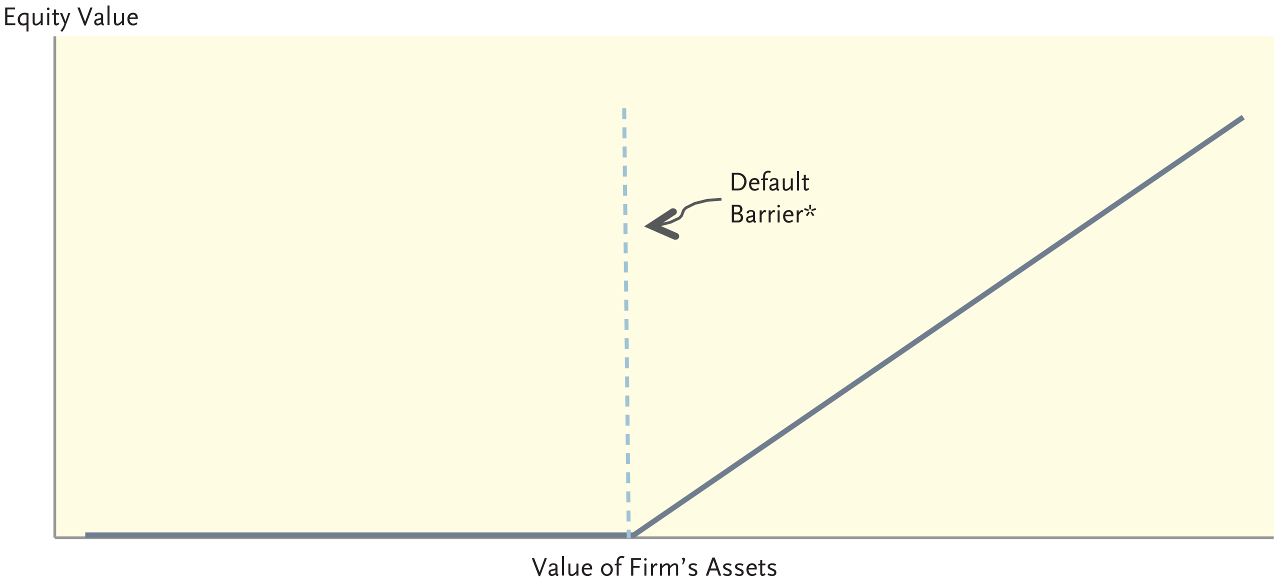

*In Merton's original model, the default barrier is equal to the value of firm's debt. Later extensions of Merton's model introduced "dynamic" default barriers based on various weights and combinations of short-term and long term liabilities, and incorporating the dynamics of risk-free interest rates and the effects of covenants.

Hence, if the face value of the firm's debt at maturity is greater than the value of the firm's assets, the equity holders are incented to default on the company's debt, allowing their "call option" to expire unexercised. Were this to occur, the bondholders become the new owners of the company's assets. Thus, conversely, owning a corporate bond is equivalent to holding a risk free bond with a short European put option on the firm's assets struck at the face value of the bond. If the value of the firm's assets is greater than the value of debt at maturity, bondholders receive their payment just like holders of risk-free debt. If the value of the assets drops below debt value, bondholders are in fact forced to "acquire" the firm's assets in lieu of the debt repayment and at a premium to the market value of the assets.

The cornerstone of this Mertonian framework is the recognition that the firm's equity holders have an implicit European call option on the assets of the firm, while the bondholders have a short position in a European put option on the assets.

Establishing the analogy between firm's equity value and a European call option allowed Merton to link the firm's asset value, debt value, and equity value through the Black-Scholes option pricing framework and, consequently, to derive the debt value of the firm. In his initial formulation, which, among other simplifications, assumed that the firm's debt is comprised of a single one-year zero-coupon bond, Merton calculated the firm's debt value based on the firm's asset value, asset value volatility, face value of debt, risk-free interest rate, and time to maturity of the firm's debt. Later, Merton modified the formula to account for coupon payments and optionality (call and put features) of the bond.

Merton further combined firm asset value, debt value, and asset volatility to derive a default probability measure. Merton calculated "distance-to-default" as the ratio of the difference of the firm's asset value and debt value to the volatility of the firm's assets[1]. Thus, distance-to-default indicates how many standard deviations the firm's asset value must drop for an event of default to occur. Alternatively, from an options-based perspective, distance-to-default may be interpreted as the number of standard deviations the equity holders' call option is in the money. As the distance-to-default decreases (as a result of elevated asset volatility, for instance), the firm's default probability goes up, increasing the likelihood that the equity holders' call option will expire unexercised.

Default Probabilities

As elegant as Merton's model is, its practical application revealed a number of imperfections. Merton's original model assumes perfect, fully-informed markets with no arbitrage opportunities, a constant risk-free rate, a single class of debt, no transaction costs or tax issues associated with bankruptcy, and default triggered only at the maturity of debt and only if the face value of debt is greater than the value of the assets. All of these assumptions were challenged in the years following the publication of Merton's pioneering work, and extensions of Merton's model were introduced to relax some of the assumptions and cope with the limitations of the model.

Notwithstanding that some of these more sophisticated models have produced superior results as compared with the original model[2], none of these modified Mertonian models have been demonstratively able to significantly improve the accuracy of the default probabilities produced by the model. Default probabilities produced by pure Mertonian models and extensions have been found to have limited predictive power because they tend to rise to alarmingly high levels only when the elevated risks are already priced in the market and it is too late for the participants to act on them. Thus, although Mertonian models do eventually produce the correct prediction, the usefulness of the information they provide is significantly limited by the delay with which these models capture the signals of elevated risks in the marketplace. Referring to this limitation of the models, some finance researchers have suggested that the predictions produced by these models are relevant only in an extremely short run and are analogous to predicting "a heart attack by observing a person dropping to the floor clutching his chest."

Perhaps the point is best illustrated by examining default probabilities generated by a Mertonian model for companies that defaulted on their debt. The chart below demonstrates one-year default probabilities for Lehman Brothers Holdings, Inc. produced by Moody's KMV[3] model.

---

[1] Needless to say, Merton's actual formula is more sophisticated and incorporates the risk-free interest rate and time to maturity of the firm's debt in addition to firm asset value, debt value, and asset value volatility.

[2] Increasingly popular "hybrid structural credit models," for example, presume that equity prices do not reflect the changes in all factors affecting firm's default probability and attempt to improve the predictive power of the pure Mertonian model by adding variables such as profitability, cash flows, liquidity, ability to refinance, etc. Some hybrid models have shown to outperform pure Mertonian models, particularly for firms in middle and low credit quality ranges.

[3] KMV was the first company to implement an extension of Merton's model to calculate default probabilities in 1989. KMV was subsequently acquired by Moody's in 2002.

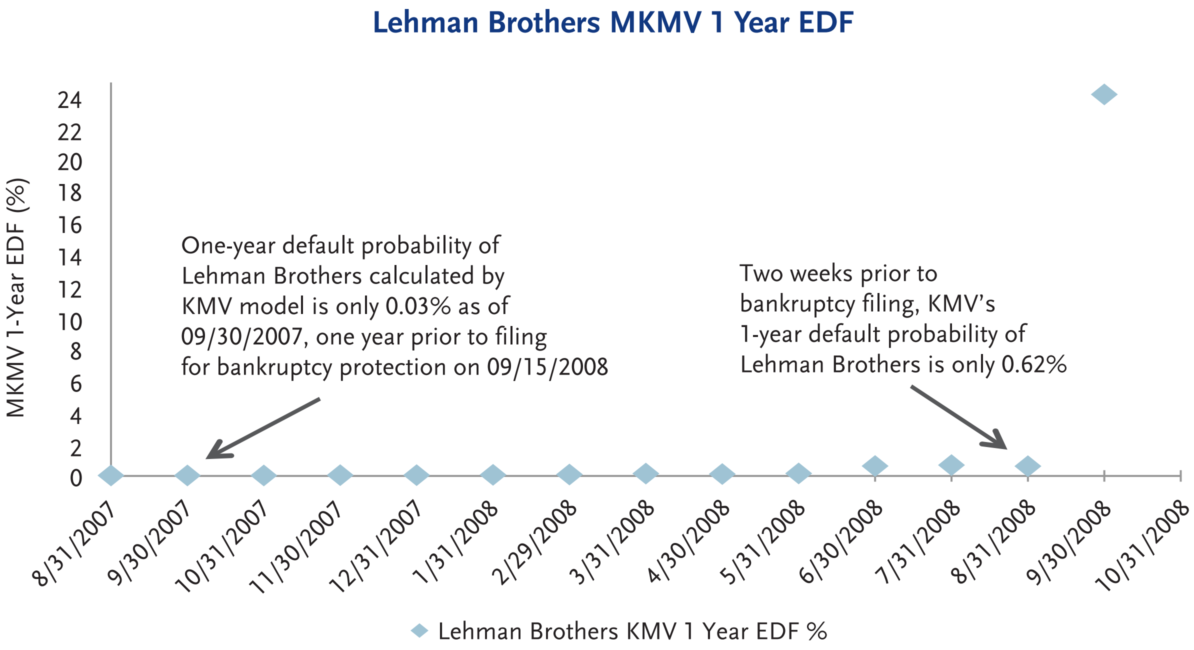

Lehman Brothers MKMV 1 Year EDF

One-year default probability of Lehman Brothers calculated by KMV model is only 0.03% as of 09/30/2007, one year prior to filing for bankruptcy protection on 09/15/2008

Two weeks prior to bankruptcy filing, KMV's 1-year default probability of Lehman Brothers is only 0.62%

Source: Moody's KMV, CreditEdge

Moody's KMV (MKMV) defines one-year Estimated Default Frequency (EDF) as the probability that a firm will default within a one-year time horizon. As shown in the chart above, one year before Lehman Brothers filed for bankruptcy protection, MKMV's model suggested a one-year default probability of mere 0.03%, and two weeks prior to bankruptcy filing, the model produced a default probability of only 0.62%. Although we do see one-year Lehman Brothers EDFs start to increase in mid-2008 from their 2007 lows, it is interesting to note that EDF decreases from 0.69% on July 31, 2008 to 0.62% on August 31, 2008. To put these EDF numbers into perspective[4] the median EDF measure of companies rated B+ by S&P was 0.83% as of August 31, 2008.

While the inability of Mertonian models to produce an early warning of default appears to be the primary concern of many end users of these models, the opposite problem, the presence of false positives (Type I errors) generated by Mertonian models, has also been widely documented in the financial literature. One of the leading companies providing risk management solutions to companies throughout the world, for example, has found that companies with more than 90% default probabilities calculated by traditional Mertonian models have an actual default rate of no more than 10%. Although generally viewed as less "harmful" than false negatives (Type II errors), false positives represent a significant limitation of structural credit models. By assigning high default probabilities to relatively low-risk securities, the models distort the risk characteristics of the securities and portfolios containing these securities, which may lead to poor risk budgeting and capital allocation decisions. Moreover, some researchers have found evidence that many structural credit models tend to artificially inflate the level of false positives (Type I errors) produced by the models in an attempt to reduce the occurrence of false negatives (Type II errors). Recognizing that the end users of these models tend to have greater aversion to false negatives than to false positives, these models exploit the trade-off between Type I and Type II errors (an increase in the level of false positives will lead to a decrease in the level of false negatives) to appear more "competitive." While manipulations of this kind may make a model more "sellable," they actually reduce the accuracy of the model and its value in the investment and risk management process.

To illustrate the illogic of this approach, imagine that we were to offer investors a model that would be absolutely certain to predict all bankruptcies, though might have some "false positives." No doubt some risk managers or investors would be interested in learning more. That is, until we told our hypothetical audience that our model "forecasts" that every firm

---

[4] In order to improve the predictive power of the model, Moody's KMV makes certain adjustments to the default barrier, firm asset volatility, and other model parameters, so the MKMV EDF numbers will differ from numbers produced by a pure Mertonian model.

will fail in the next twelve months! Yes, the model is perfect in that would miss no bankruptcy, but the model is also perfectly useless. Needless to say, a useful model must be able to reliably discriminate between situations where bankruptcies are likely versus those where they are not.

Credit Spread Puzzle

The other major limitation of Mertonian models observed by structural credit model practitioners is that credit spreads predicted by these models are systematically lower than the actual spreads observed in the marketplace. This phenomenon, dubbed the "credit spread puzzle," is believed to be grounded in the simplifying assumptions of Merton's model.

Some finance researchers have tried to explain the "credit spread puzzle" by questioning the use of the U.S. Treasury curve as the risk-free curve for calculating corporate spreads. They contend that because of the extreme liquidity and special status of U.S. Treasury securities, the U.S. Treasury rates are lower than the true risk-free rate. Therefore, credit spreads observed in the marketplace reflect not only the compensation demanded by investors for just corporate credit risk but also a liquidity premium. Recent studies have confirmed this intuition and have shown that liquidity risk is priced in credit spreads. Accordingly, some of the more advanced models, including MKMV, adjust credit spreads to account for liquidity premium by estimating an implicit "corporate risk-free curve." While these modified models show an improvement over their unadjusted peers, they are unable to "solve" the credit spread puzzle. Moreover, it has been observed that some of the newer, more sophisticated models tend to underpredict spreads for safer bonds, while significantly overstating the spreads of riskier firms. Examination of implied credit default swap (CDS) spreads produced by MKMV's model reveals that the model tends to underpredict spreads for shorter maturities, while overpredicting spreads for longer maturities. The chart below shows the difference of 1-year EDF-implied CDS (EICDS) spreads[3] generated by the MKMV model and actual 1-year CDS spreads observed in the marketplace (market spreads obtained from Barclays Capital) for six large, fairly liquid firms representing various sectors of the economy. The graph clearly illustrates the credit spread puzzle, consistent with the results documented in financial literature. Moreover, we observe that the underprediction becomes more severe during periods of high volatility, such as the second half of 2008 and the first half of 2009.

1Year MKMV EICDS Spreads[5] – 1 Year Market CDS Spreads[6]

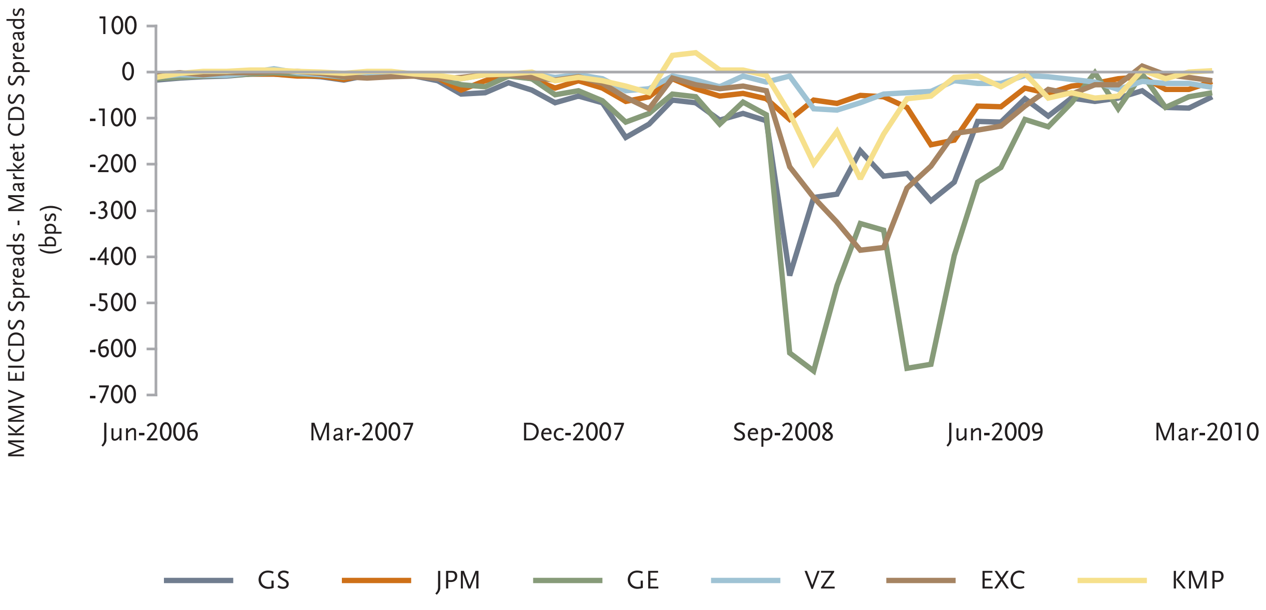

---

[5] The 1 Year EICDS Spread is the spread on a one-year CDS for a given entity, as calculated by the Moody's KMV valuation model.
[6] Market CDS spreads are estimated by spreads obtained from Barclays Capital.

Interestingly, implied CDS spreads produced by the MKMV model for longer maturities (e.g. 7-year spreads and 10-year spreads) are consistently higher than the actual spreads observed in the market. The chart below illustrates the differences of 10-year EDF-implied CDS (EICDS) spreads produced by Moody's KMV model and actual 10-year CDS spreads observed in the market for the same six firms for which MKMV was shown to underpredict 1-year spreads. The pattern of increasing deviations from actual spreads in the times of unusually high market volatility is also observed here, although it is not as pronounced as in the case of 1-year spreads.

10-Year MKMV EICDS Spreads[7] – 10-Year Market CDS Spreads[8]

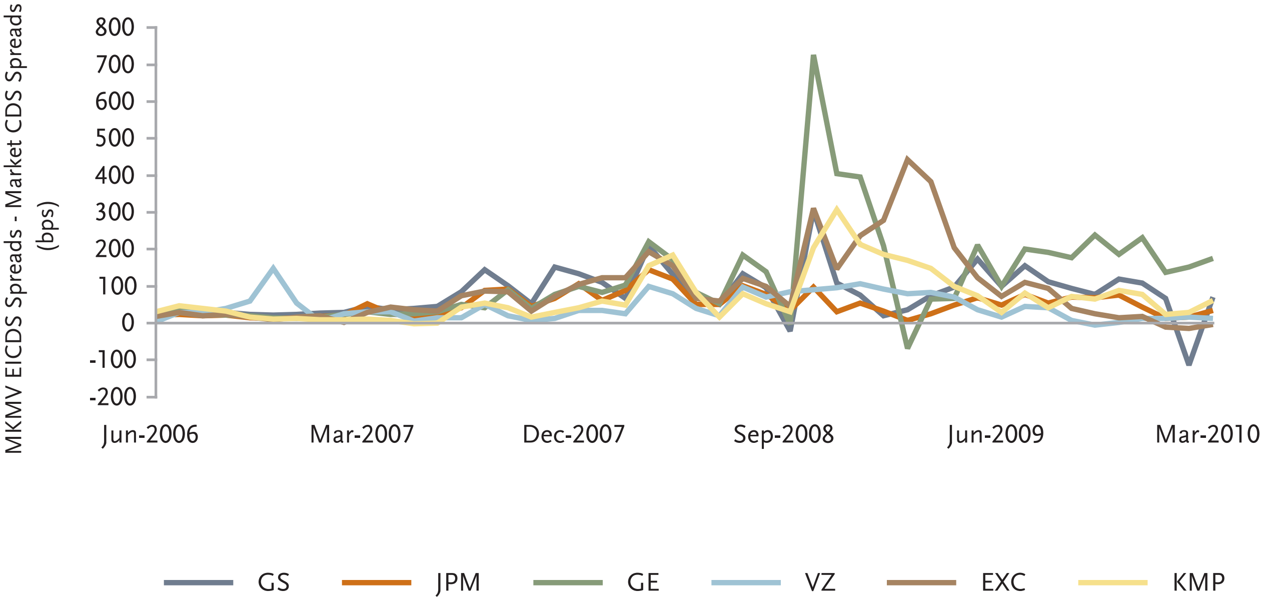

Another potential contributor to the credit spread puzzle identified in the financial literature is the "information risk." The key insight here is that the values of certain inputs of structural credit models (e.g. firm asset value, volatility of firm's asset values) cannot be observed directly, but rather are estimated based on available market data, such as equity prices. Therefore, higher actual spreads observed in the market may reflect the additional premium required by investors to compensate for "information risk," i.e. risk associated with estimating unobservable model parameters. While the intuition behind this theory is sound and the theory may potentially offer a rationale for under-predicting the spreads, it seems to fall short of solving the credit spread puzzle.

Mertonian Models vs. "Reduced Form" Models

There is hardly a single assumption in the original Mertonian framework that has not been challenged by finance researchers and model practitioners in the last 35 years. Recognizing the relatively poor performance of Mertonian models, finance researchers approached the task at hand with a great deal of enthusiasm and introduced dozens of "new," more sophisticated credit default models. While most of these "new" models are to a certain degree modifications of Merton's model, some researchers questioned the capability of Mertonian (structural credit) models to produce accurate default probabilities and introduced a new framework for calculating probability of corporate debt default, called the "reduced-form" framework.

[7] The 10 Year EICDS Spread is the spread on a ten-year CDS for a given entity, as calculated by the Moody's KMV valuation model.
[8] Market CDS spreads are estimated by spreads obtained from Barclays Capital.

In Merton's original framework, default is triggered only at bond maturity if the value of firm's assets is less than the face value of maturing debt. Extensions of Merton's model later introduced "default barriers" – default-triggering levels of a firm's asset value which allow for default to be triggered prior to the maturity of the bond. However, even the most sophisticated Mertonian models have the firm's debt level as the sole determinant of default.

The new, reduced-form models, introduced in the 1990s, depart from the pivotal assumption embedded in Mertonian models that default may be triggered only by the inability of the firm to repay its debt. Rather, the reduced-form approach assumes that default may be triggered by various other factors, in addition to the inability to repay debt, and that it may be triggered at any point in time. For instance, a firm with no outstanding debt may file for bankruptcy protection because it does not have sufficient cash (or other liquid assets) to meet its current liabilities, such as lease or salary payments. Additionally, firms with "healthy" balance sheets have in the past been forced to file for bankruptcy protection because of unforeseen legal liabilities.[9] Distressed debt investors are also well aware of situations in which one class of bondholder (e.g. senior secured lenders) might find bankruptcy or liquidation preferable to having management continuing to operate the business. Finally, companies may also engage in coercive debt exchanges that technically circumvent bankruptcy, but are, in actuality, a default by another name.

While structural credit models assume that at each point in time market participants know the value of the firm's assets and can predict the occurrence of default, the reduced-form approach assumes that a firm's default is unpredictable and is driven by a default intensity, an exogenous process which is a function of variables available to the investors. Yet, practitioners are well-versed in understanding that companies often have "hidden" liabilities or liabilities that are at least very difficult to quantify (indeed, liabilities have a way of "growing" in bankruptcy, suggesting that the notion of a liability isn't so much a hard and fast accounting concept as it is a subjective claim predicated on a legal theory).

Reduced-form models choose inputs from a broad array of firm-specific as well as macroeconomic variables based on their explanatory power in fitting historical default curves. The model then calculates default probability based on historical probability distributions and correlations between the variables selected. Although reduced-form models are generally credited with producing more accurate results compared to Mertonian (structural credit) models, the advocates of Mertonian models argue that the fundamental weakness of reduced-form models is the lack of theoretical link between the financial situation (asset-liability structure) and the credit risk of the firm. A common retort is that the real world problem of forecasting defaults has shown itself to be much more complex than the Mertonian framework makes allowance for.

Conclusion

Merton's pioneering work published in 1974 sparked an explosion of research on corporate debt valuation and default probability. This resulted in scores of research publications and dozens of new models, which were mainly introduced in an attempt to improve the overall performance of Mertonian models. Moreover, some researchers challenged the ability of Mertonian models to produce consistently accurate results and introduced an alternative ("reduced-form") framework for evaluating credit default probability. Some of the new credit default models introduced after the publication of Merton's pioneering work were shown to improve the accuracy of the model, and the commercial implementations of these models became popular tools used by credit investors and risk managers. Despite the great strides made by finance researchers in the last four decades, however, even the most advanced credit default models still have serious limitations and biases, which must be understood and carefully considered in the context of investment decision making and risk-management processes.

[9] The recent oil spill in the Gulf of Mexico offers a fresh example of a financially healthy, growing firm being "shocked" by an unexpected event and the unexpected liabilities associated with it. Although BP has not filed for bankruptcy protection and is trying to meet its liabilities with cash raised through asset sales, the mounting liabilities from the spill have made the markets extremely nervous, sending the one-year CDS spreads on BP from 25 basis points in April, 2010 to north of 900 bps in June, 2010.