Better Ingredients

Christopher C. Finger
chris.finger@riskmetrics.com

Robert Stamicar
robert.stamicar@riskmetrics.com

April 2005

> *…no cook, however creative and capable, can produce a dish of a quality any higher than that of the raw ingredients.*
> Alice Waters

Chez Panisse in Berkeley, California has, over the last thirty years, become one of the most decorated restaurants in the United States. The most cited reason for the restaurant's renown is not the complexity of the dishes, but rather the commitment of executive chef Alice Waters to obtaining the highest quality, seasonal, local ingredients.

Stretched as the analogy might be, there is a lesson from Chef Waters we can apply to modeling. CreditGrades™ is our implementation of the structural model for credit risk. Simply put, the model takes information about the value, volatility and leverage of a firm, and predicts the spread at which the firm's credit should trade. In general, the model performs well, but there is much we can learn from the cases where it does not. Though it is natural to approach modeling deficiencies by tinkering with the model, it is often more beneficial to revisit the data we feed in. In fact, the structural model is particularly amenable to this approach, since it links so many elements.

In this note, we will examine the structural model's performance with two firms, Vivendi and General Motors, that either were or are significant credit stories, and for which the most straightforward model implementation appears flawed. Though there are plenty of avenues we could take to enhance the model itself, we will focus instead on the model inputs. Our approach, then, is not to make our ragout more complex, but instead to make it with fresher mushrooms.

The standard framework

The structural framework for modeling credit risk links a firm's ability to pay on its liabilities to the level of its current assets and the evolution of those assets through time. In very simple terms, we specify the level of the assets, the volatility of the assets and the level of the liabilities, and compute the likelihood that the value of the assets will fall below the liability level over the time period in question. Additionally, at least at a point where we liquidate the firm, the value of the firm's equity would be simply the excess of the value of the assets over the liabilities. Thus, the evolution of the firm's assets drives both the firm's equity and credit, and therefore, the model framework creates a link between the equity and credit markets. In practice, we may thus choose the market we believe to have richer information, and apply that information to the other market.

Distinct implementations of the structural framework vary in precisely how they apply the intuition above to calculate the equity value from the asset parameters. In a more robust implementation, we observe that the equity value has an option-like payoff in liquidation: it is the excess of assets over liabilities, but cannot be worth less than zero. The equity value today is the discounted expected value of this future payoff, and is governed by an option pricing formula. In the CreditGrades implementation, we apply the approximation that the equity value is equal to the excess of assets over liabilities, not just at the time of liquidation, but at any time, and arrive at a much simpler relationship between assets and equity.[1] In the spirit of our introduction, it will not be our focus to compare implementations, but rather to examine the data that is common to any implementation.

Applications of the structural framework have evolved as various markets have gained in liquidity. Historically, the most common applications of the structural framework relied on the equity market as the source of information. Further, such applications typically focused on lending decisions or warning signals, and therefore traditionally sought to provide accurate default probabilities (in a sense, better ratings) rather than true pricing information. This made sense, as pricing information in credit was scarce, making it difficult either to calibrate or apply a model to true credit prices.

A change, not so much to the model framework as to its application, came about in the early part of this decade. With the credit derivatives market having matured, market practitioners recognized that the structural framework could be utilized beyond traditional credit applications, and provide hedging, relative value and risk information to a trading operation. Thus, with CreditGrades, the goal has still been to look ultimately to the equity market as the source for information, but to apply the structural model on the equity information in order to produce indicative prices for credit. This has changed how we evaluate structural models: we now require that a model provide indications of credits that appear to be over- or undervalued relative to their peers, and to provide timely indications of changes in the price of a firm's credit.

As the credit market continues to grow, it is no longer obvious that the equity market should be the first place we look for information. Rather, there are times where a firm's credit will lead significant equity changes, or will remain stable in the face of an equity sell-off, indicating that there is no fundamental problem with the firm. Further, as we will see, equity options can add significant information to the model, and so it is desirable to shop for ingredients at this market as well.

Incorporating options

There are no additional model assumptions required to price options within the structural framework. As we explained before, simply to value equity and credit, we must specify a model for the firm's assets. Since we have specified how the assets evolve, and the model gives us the value of the equity as a function of the assets, we have also (implicitly) specified how the equity evolves. This is sufficient to derive the distribution of future values of the equity, and to value options on the equity. Of course, distinct model implementations imply distinct equity

[1] In fact, the approximation is that the equity is equal to the asset value less the value of the assumed recovery value of the liabilities if a default were to occur.

processes and option prices, but the model framework is rich enough to price equity options in all cases.

One attractive feature of the structural model as applied to equity options is that the model implies a skew in option prices. Typically, the evolution of the assets is described by a process with constant volatility. Thus, it is equally likely that the assets move by 10% if their current level is 100 as if their current level is 1000. On the other hand, if the liability level is 89, the impact of a 10% change in asset value will have a much greater effect on the equity value if the assets are currently at 100 than if they are at 1000. This contrasts with the Black-Scholes framework, in which the equity volatility is assumed constant. The implication is that equity becomes more volatile as its value falls and the firm's leverage increases. As a result, options with lower strike prices are more valuable than a standard Black-Scholes model would predict, as is the typical case in the marketplace.

With options incorporated into the framework, we may think of the model as a way to connect fundamental firm characteristics (asset value, asset volatility, liability, and recovery rate) to prices on a variety of contracts (equity, equity options at various strike prices, and credit). This wide array of inputs and outputs affords us numerous ways to apply the model. At one end of the spectrum is the fundamental approach: estimate the model input parameters from first principles, and then calculate the model prices of the various financial instruments. At the other end of the spectrum lies the market-based approach: observe the market prices of the financial instruments, and solve for the model parameters that allow us to recover these prices.

In practice, we will typically blend the fundamental and market-based approaches, depending on our relative confidence in the observed market prices and our parameter estimations. For our two examples, we will examine the CreditGrades implementation. We will utilize five-year Credit Default Swaps (CDS) and one-year At-the-Money (ATM) equity options. Following the standard CreditGrades approach, we describe liability levels on a normalized basis, using the quantity Debt-per-Share (DPS).

Vivendi

Our first example is Vivendi Universal, the French media and telecommunications conglomerate. Though a relatively quiet credit today, with an investment grade rating and default swap spreads comfortably under 50bp, Vivendi was one of the biggest headlines of 2002, when a cash crisis and allegations of accounting improprieties drove it to the brink of bankruptcy.

We compare two implementations of the structural model for Vivendi. In the first approach, *CG Fundamental*, we observe the firm's equity price, and estimate the equity volatility using historical returns. We assume a recovery rate, and estimate DPS from the firm's balance sheet. Using the CreditGrades model, we derive an asset volatility, and then calculate the fair value of the CDS. In the second approach, *CG ATM Vol*, we observe both the firm's equity price and the price of an ATM equity option. We assume a recovery rate and estimate DPS as before, then infer the asset volatility that correctly prices the option and then calculate the fair CDS level. We plot the CDS and both model spreads in Figure 1.

From the beginning of 2002, Vivendi's stock started to drop, falling about 40% by late April, though its

Figure 1: Vivendi. Credit Default Swap and CreditGrades spreads

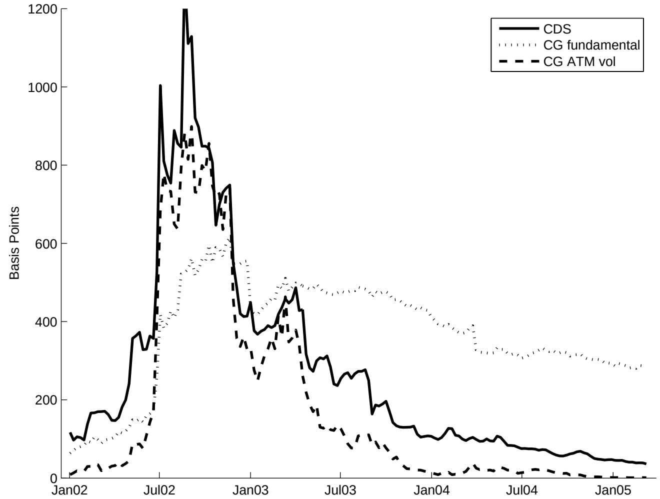

CDS spreads remained under 200bp. In early May, concerns over the pace of Vivendi's acquisitions and its cash positions arose, the company announced a writedown of EUR 15 billion, the stock dropped more quickly, and the CDS spread moved to 400bp in just two weeks. The CG Fundamental model provided relatively accurate spreads at the beginning of the year, and its spreads did widen in sympathy with the stock depreciation, but the model's spread widening significantly lagged the actual widening.

On July 2, 2002, French newspaper *Le Monde* published allegations about Vivendi's accounting practices. Vivendi stock plunged as much as 40%, and Chief Executive Jean-Marie Messier resigned that evening. Vivendi's long-term debt was downgraded to junk status, and spreads on Vivendi CDS widened as much as 500bp in a single day. The CG Fundamental model did predict some spread widening, consistent with the sharp drop in stock price, but this widening was not nearly as severe as the actual spread move. In fact, while CDS spreads bounced between 800 and 1200bp for three months, the CG Fundamental crept from 400 to 600bp.

The crisis lasted until late October, by which time new Chief Executive Jean-Rene Fourtou had taken over, and Vivendi had begun the process of divesting of many of its previous acquisitions. The stock price has recovered, and is now near its May 2002 level. Vivendi debt is again investment grade; its CDS spread has been below 50bp for all of 2005, and continues to tighten. In contrast, the CG Fundamental spreads have tightened very slowly, and are currently just below 300bp.

While examples exist of firms that ran into crisis and whose spread was well predicted by the CG Fundamental model, the model's performance with Vivendi is disappointing. On the day the accounting allegations emerged, when the CDS spread widened by 450bp, the CG Fundamental spread only widened by 88bp. This is an indication that the market's credit concerns were not fully manifested in the stock price level. Rather, much of the concern was due to the uncertainty surrounding the allegations. Similarly, once the crisis had passed, the stock price recovered, and the uncertainty subsided as well. These changes in uncertainty seem to be present in the CDS levels, and are clearly present in prices of options on Vivendi stock: implied volatility for at-the-money options was at 40% in May 2002, rose to around 80% during the crisis and is now at 25%.

Intuitively, we expect that the dynamics of the equity options should help with forecasts of the spread. Indeed, spreads predicted by the CG ATM Vol model spiked precisely when the crisis began, remained at crisis levels through October, and then subsided. In fact, while during the crisis, the model produced comparable spread levels to the market, the model spreads tightened more quickly than the market after the crisis, and are now only a few basis points. Thus, the credit, equity and options markets for Vivendi appear to have traded in perfect harmony during the crisis, while today it appears that credit is cheap[2] relative to where the equity and options are trading.

Hull et al (2005) take a step further away from the fundamental model, calibrating both an asset volatility and an implied liability level in order to match both an at-the-money and an out-of-the-money option.[3] They find that the model performs reasonably well versus the analog of our CG Fundamental model in distinguishing between low and high spread names. However, in a simple regression exercise, they observe that the option skew does not always add information to the at-the-money volatility for explaining spreads. Furthermore, they comment that the skews they observe are often significantly greater than those that are typically obtained through the structural model.

We apply a similar approach to Vivendi. We fit an implied DPS to two option prices and calculate a model spread, expecting that we might bring the model spreads in the non-crisis periods closer to the actual market spread. This method does provide marginally closer spreads for the recent two months, but in general, it overcorrects, producing unreasonably high spreads. Thus, our experience in this case is consistent with Hull et al's observations: the skew adds only marginal information, and is often greater than the typical range predicted by the structural model.

General Motors

While Vivendi is a nice name to write about from a historical perspective, particularly since things are quiet today, General Motors (GM) is arguably the biggest current story in the credit markets. As the largest corporate bond issuer in North America, GM is always near the center of attention, and its significant pension liabilities and struggling sales have garnered it even more scrutiny in the last two years. On March 16, GM issued a significant profit warning and Standard and Poor's revised its GM outlook; GM stock fell by almost 15% and its CDS spread widened by 80bp.

From the perspective of the structural model, GM has always been difficult to tackle. As with all names, estimating an empirical equity volatility is challenging. With GM, however, there is the added complication of capturing its true leverage. Over 80% of outstanding GM bonds are issued by General Motors Acceptance Corporation (GMAC), its financial services subsidiary. Because GMAC operates more like a bank than an industrial corporation, much of its debt is secured, and it is difficult to argue that all of the GMAC debt contributes to the overall leverage of GM. It is difficult to make any conclusion on the

[2] That is, spreads are wide.

[3] Hull et al use the implementation where equity is priced formally as an option on firm assets.

Figure 2: General Motors. Implied Debt-per-Share and 50-day moving average

Figure 3: General Motors. Credit Default Swap and CreditGrades spreads

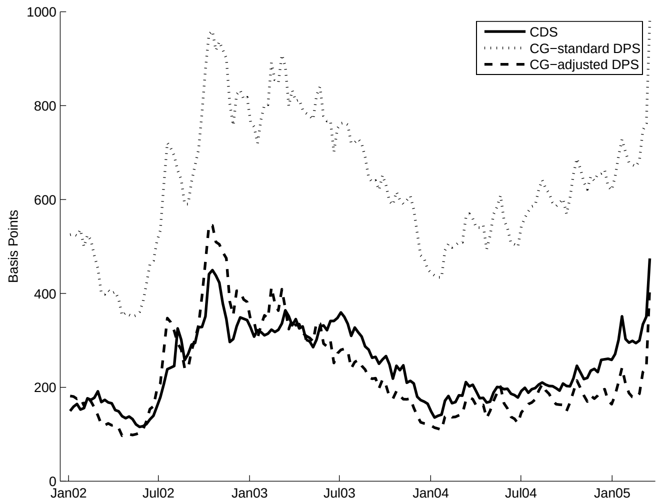

liabilities, however, when we must simultaneously decide on our volatility estimate, and we only have a single market observation (the CDS spread) to guide us.

Our lesson from the Vivendi example is quite useful here. Rather than relying on a historical estimate for equity volatility, we may confidently infer the model asset volatility from the ATM equity option on GM. Furthermore, since GM credit spreads are among the most liquid in the market but we are unsure of how to characterize its leverage, we treat the spread as an input and infer the value for Debt-per-Share such that the model recovers the observed spread. This approach is similar to the approach with the option skews discussed above, where we infer asset volatility and DPS from the two option prices. For the reasons discussed there, we are less confident with DPS inferred this way, and opt instead to calibrate to the ATM option and CDS spread.

We plot the historical implied DPS in Figure 2. There are two key observations from this figure: first, just prior to July 2003, there was a sharp rise in the DPS, from about USD 120 to USD 160; second, the implied DPS has been quite stable since July 2003, even including recent events.

The sharp rise in DPS is coincident with GM's June 2003 USD 13B bond issuance. This issuance, intended primarily to fund GM's pension shortfall, was unusual for GM in that most of the issuance (USD 10B) was issued by General Motors Corporation, not the GMAC subsidiary. The June 2003 issuance accounts for about USD 20 of the GM DPS. This is less that the USD 40 increase we see in the implied DPS, but it is at least roughly the magnitude of the actual increase. This, coupled with the timing of the increase, is evidence that the implied DPS is a meaningful quantity, and not simply a fudge factor to correct for poor model performance.

The recent stability of the implied DPS implies that the CG ATM Vol model could describe spread moves quite well if only we could arrive at a reasonable liability level. Since July 2003, the level of the implied DPS has been roughly 20 to 25% of the overall GM DPS.[4] At present, the GMAC subsidiary accounts for approximately 85% of outstanding GM bonds. Returning to the intuition that the effective GM DPS should include the GM debt, plus some small portion of the GMAC debt, it is sensible to use a model DPS equal to 20% of the overall value.

We apply the CG ATM Vol model using the adjusted GM DPS. We plot the results, along with the CDS spread and the unadjusted model, in Figure 3. The results are very encouraging. Even in 2002, when the implied DPS was significantly different from the level we set, the model spread tracks the market tightly. Further, the model picks up even the most recent spread widening, suggesting that the relationship between equity, options and credit continue to hold despite significant upheaval with GM.

Applications

For both of our examples, better ingredients do improve our dish. By inferring volatility from equity options and adjusting our DPS estimate, we see that the spread predicted by the structural model tracks very closely to the actual credit spread. The implications of this for trading are more accurate relative value signals. What is more, trades based on an indication that credit is mispriced relative to equity and

[4] That is, the DPS calculated using all liabilities of GM and its subsidiaries.

options will consist of credit, equity and option positions.

But how does our risk management benefit from the enhanced model accuracy? For firms whose credit trades liquidly, we may assess spread volatility directly from the market, but the model affords us a better understanding of what drives credit spreads. Not only is there a component of spread volatility that is linked to equity prices, but also a component driven by equity option volatility. Thus, exposures to credit spreads contribute to our overall exposure to the equity markets, as well as our exposure to equity volatilities. In a sense, we have established that there is vega risk in the simplest of credit positions, and a solid understanding of our risk exposures should incorporate this.

For firms whose credit does not trade liquidly, the enhanced model opens up more possibilities. Beyond more information about risk exposures, the model gives us a realistic view of spread volatility itself. Extending the ideas in Mina and Ta (2002), if we can forecast the future volatility of both equity and implied volatility, we can apply the model to assess the Value-at-Risk of a position or portfolio due to spread movements. Finally, the model provides for more robust warning signals, providing more accurate, and less obvious, signals than a model that refers to equity price alone.

Further reading

- Finger, C. (ed.) (2002). *CreditGrades™ technical document*, RiskMetrics Group.
- Hull, J., Nelken, I., and White, A. (2005). Merton's model, credit risk and volatility skews, *Journal of Credit Risk*, **1**(1): 3–28.
- Merton, R. (1974). On the pricing of corporate debt: the risk structure of interest rates, *Journal of Finance*, **29**(2): 449–470.
- Mina, J. and Ta, T. (2002). Estimating issuer-specific risk for corporate bonds, *RiskMetrics Journal*, **3**(1): 7–26.