



IMF Working Paper

Probabilities of Default and the Market Price of Risk in a Distressed Economy

Raphael Espinoza and Miguel Segoviano

IMF Working Paper

Strategy, Policy and Review Department

Probabilities of Default and the Market Price of Risk in a Distressed Economy

Prepared by Raphael Espinoza and Miguel Segoviano¹

Authorized for distribution by Catherine Pattillo

April 2011

This Working Paper should not be reported as representing the views of the IMF.

The views expressed in this Working Paper are those of the author(s) and do not necessarily represent those of the IMF or IMF policy. Working Papers describe research in progress by the author(s) and are published to elicit comments and to further debate.

Abstract

We propose an original method to estimate the market price of risk under stress, which is needed to correct for risk aversion the CDS-implied probabilities of distress. The method is based, for simplicity, on a one-factor asset pricing model. The market price of risk under stress (the expectation of the market price of risk, conditional on it exceeding a certain threshold) is computed from the price of risk (which is the variance of the market price of risk) and the discount factor (which is the inverse of the expected market price of risk). The threshold is endogenously determined so that the probability of the price of risk exceeding it is also the probability of distress of the asset. The price of risk can be estimated via different methods, for instance derived from the VIX or from the factors in a Fama-MacBeth regression.

JEL Classification Numbers: G13; G21

Keywords: Price of risk; credit default swap; CDS; risk-neutral probability

Author's E-Mail Address: respinoza@imf.org; msegoviano@cnbv.gob.mx

¹ Miguel Segoviano is Director General of Risk Analysis and Quantitative Methodologies at the Mexican Financial Authority (Comisión Nacional Bancaria y de Valores). The views expressed in this Working Paper are those of the authors and do not necessarily represent those of the IMF or of the Mexican Financial Authority. We are grateful to Thanawan Chaiwatana for outstanding research assistance. This paper benefitted greatly from discussions with Tobias Adrian, Carlos Caceres, Vincenzo Guizzo, Huyn Shyn, Dimitrios Tsomocos, and from participants at the IMF and at the LSE/CCBS Bank of England conference on Measuring Systemic Risk and Issues for Macroprudential Regulation. Any errors are solely the authors' responsibility.

Contents	Page
I. Introduction	3
II. Theoretical Background	4
A. Risk-neutral Probabilities of Default and the Market Price of Risk	4
B. Conditional Expectation of the Market Price of Risk	6
C. Estimating the Price of Risk.....	7
III. Endogenous Threshold.....	9
IV. Application to US Banks during the Crisis.....	11
V. Conclusion	12
References.....	14
 Figures	
Figure 1. Ratio of CDS-implied Probability of Distress to Moody's KMV EDF	4
Figure 2. Uniqueness of the Threshold	10
Figure 3. Adjustment Factor and Market Price of Risk	11
Figure 4. Estimated Probabilities of Default.....	13

I. INTRODUCTION

During the credit crisis of 2007-2009, the estimation of default probabilities of banks has been a focal point of interest. Default probabilities can be estimated using markets' assessment, as captured by CDS spreads, or using models of the value of the firm derived from the Black-Scholes-Merton model (Black and Scholes, 1973; Merton, 1974). The latter method has been popularized by the Moody's KMV Estimated Default Frequency (EDF) dataset, which provides default probabilities for many of the largest companies in the world, and most U.S. banks. CDS spreads are also widely used, although they are only proxies for default probabilities as they are influenced by the market price of risk (the cost of insurance) as well as the probabilities of distress of the firms.

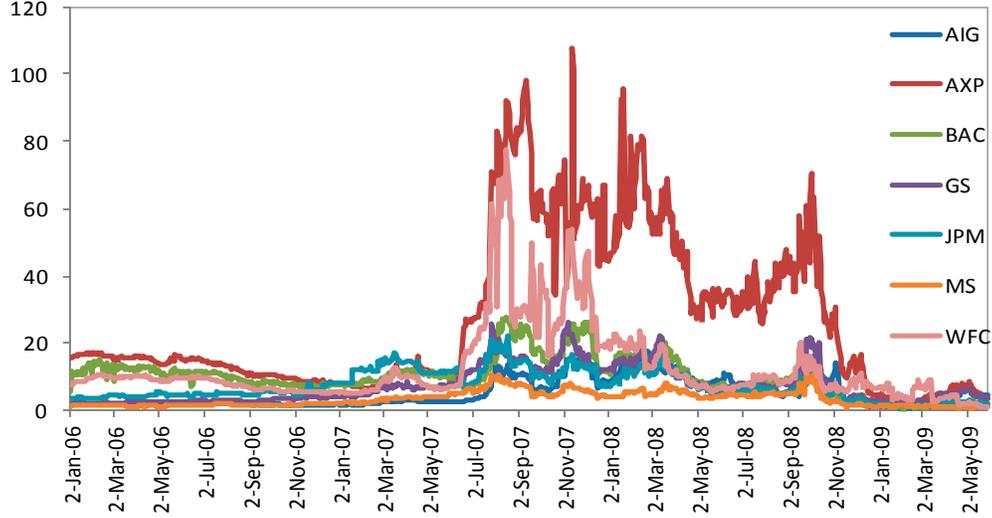
The two measures of probabilities of stress have diverged markedly during the post-Lehman times. We show in Figure 1 the ratio of the CDS-implied probability of default over the EDF probability of default for different U.S. banks. This ratio, which could be interpreted as the market price of insurance if the EDF and the CDS spread were representing perfectly the probability of default and the risk-neutral probability of default, has also varied dramatically *across banks*. It has been shown to vary across sectors (Berndt et al., 2005), and to correlate negatively with real activity (as one would expect from a pricing kernel—see section II) and positively with real interest rates (Amato, 2005, Amato and Luisi, 2006).

The objective of this paper is to provide a method for computing the market price of risk under distress and therefore the probabilities of default implied by the CDS spreads. The objective is to provide an alternative measure to the EDFs, which are based on slowly-moving balance sheet information and have been lagging market information during the 2007-2009 crisis. The probabilities of distress in the recent abnormal times cannot be estimated using past data, so the method proposed by Jackwerth (2000) to compute risk aversion and actual probabilities is also ill-suited.

In our paper, the calculations assume a one factor model (for instance a Consumption CAPM, although this is not needed for the results) and show how information on the mean and the variance of the factor can be used to derive the market price of risk under a situation of distress. This market price of risk is the conditional expectation of the price of risk, conditioning on it exceeding a threshold (section II).

The factor we use is a transformation of the VIX (which has been shown to correlate strongly with the first principal component of asset returns). The approach is free of any assumption on the shape of the utility function (e.g., the coefficient of risk aversion) and of assumptions on stationarity of returns, but nonetheless it allows us to strip the effect of risk aversion on asset prices. However, the two fundamental assumptions are that pricing is based on a one-factor model, and that this factor is normally distributed.

Figure 1. Ratio of CDS-implied Probability of Distress to Moody's KMV EDF



Source: Moody's KMV and Bloomberg

We finally show that the market price of risk under distress has to be calculated jointly with the probability of distress to ensure the conditioning threshold is compatible with the probability of default (section III). We apply our method to the major U.S. banks during the crisis (section IV) and show that CDS-implied probabilities of default overestimated credit risk by roughly 50 percent. Section V concludes.

II. THEORETICAL BACKGROUND

A. Risk-neutral Probabilities of Default and the Market Price of Risk

In this section we review the basic asset pricing framework linking the price of risk to the CDS-implied probabilities of default. A one factor model can be derived from a consumption Euler equation:

$$P_t = E_t \left[\beta \frac{u'(c_{t+1})}{u'(c_t)} x_{t+1} \right] \quad (1)$$

where P_t is asset price at time t and x_{t+1} is the one period-ahead payoff. The stochastic discount factor is defined as

$$m_{t+1} = \beta \frac{u'(c_{t+1})}{u'(c_t)} \quad (2)$$

and the linear pricing formula for the one-factor model is therefore:

$$P_t = E_t [m_{t+1} x_{t+1}]. \quad (3)$$

Writing the gross return $R_{t+1} = \frac{x_{t+1}}{P_t}$, the pricing formula is equivalent to

$$1 = E_t[m_{t+1}R_{t+1}]. \quad (4)$$

Applying the previous equation to a risk-free asset:

$$1 = E_t[m_{t+1}R_{t+1}^f] \text{ or } R_{t+1}^f = \frac{1}{E_t[m_{t+1}]} \quad (5)$$

If the states of nature at time $t+1$ are indexed by s , and the probability of state s is $\pi(s)$

$$P_t = \sum_s \pi(s)m(s)x(s) = E_t[m_{t+1}x_{t+1}] \quad (6)$$

The risk-neutral probability $\hat{\pi}$ is such that²

$$\hat{\pi}(s) \equiv R_{t+1}^f m(s)\pi(s) \quad (7)$$

This probability measure does not correspond to any actual probability of nature-in particular it does not imply that the asset pricing model is based on risk-neutral investors. Nonetheless, the term ‘‘risk-neutral probability’’ is used in the literature because this is the probability measure that a risk-neutral investor would need to ‘‘believe in’’ to agree on the asset prices given by the markets. Indeed, the price of an asset can be written as

$$P_t = \frac{1}{R_{t+1}^f} \sum_s \hat{\pi}(s)x(s) = \frac{\hat{E}(x)}{R_{t+1}^f} \quad (8)$$

where $\hat{E}(x)$ is the expected value of x_{t+1} associated with the risk-neutral probability measure $\hat{\pi}$. For a risk-neutral investor that believes that $\hat{\pi}(s)$ are the actual probabilities of nature, the price P_t given by the markets is coherent with the payoff x_{t+1} .

The risk-neutral probability $\hat{\pi}(s)$ is what is obtained from CDS spreads. Indeed, $\hat{\pi}(s) / R_{t+1}^f$ is the price of an asset that pays $x(s) = 1$ dollar in the state of distress s (from equation 8). A common approximation is $\hat{\pi} = \frac{S_N}{(1-K)}$ where S_N is the CDS spread of bank N and K the recovery rate (assumed to be at 60 percent), and R_{t+1}^f can be proxied by the OIS rate. The missing element is therefore the market price of risk $m(s)$.

² $\hat{\pi}$ is a probability measure since all $\hat{\pi}(s)$ are positive (from the fundamental theorem of finance, $m(s)$ is positive in absence of arbitrage opportunities). Furthermore, $\sum_s \hat{\pi}_d(s) = \sum_s R^f m(s)\pi_d(s) = E_t[R^f m_{t+1}] = 1$

B. Conditional Expectation of the Market Price of Risk

Since our focus is on the state of nature where the banks are under distress, we need not estimate $m(s)$ for all states of nature. We group the states of nature under two headings: *distress*, and *no distress* and rewrite the linear pricing formula:

$$\begin{aligned} P_t &= E_t[m_{t+1}y_{t+1}] \\ &= \pi_d E_t[m_{t+1}y_{t+1} \mid \text{distress}] + (1 - \pi_d) E_t[m_{t+1}y_{t+1} \mid \text{no distress}] \end{aligned}$$

The relationship between the risk-neutral probability under stress and the real probability of stress is given by equation (7)

$$\pi_d = \frac{\hat{\pi}_d}{R^f E_t[m_{t+1} \mid \text{distress}]}$$

The market price of risk under stress $E_t[m_{t+1} \mid \text{distress}]$ is not observable and our objective is to estimate it from market prices. Calibrations based on the shape of the utility functions, derived from equation (2), could also be used, but the link between asset prices and utility functions is subject to many difficulties, as evidenced by the numerous puzzles spurred by the consumption-CAPM literature.

Our method relies on the estimation of $E_t[m_{t+1} \mid \text{distress}]$ based on the mean and the variance of m_{t+1} , which can both be measured.

Indeed, $E_t[m_{t+1}]$ is deduced from the Overnight Indexed Swap (OIS) rate, following equation (4). Furthermore, the volatility of m_{t+1} ($\text{Var}_t[m_{t+1}] / E_t[m_{t+1}]$) is the *price of risk*, which has been the subject of much attention in the asset pricing literature. Finally distress is defined as the situation where the market price of risk exceeds a certain threshold (see below).

Under the additional assumption that m_{t+1} is normally distributed, $E_t[m_{t+1} \mid \text{distress}]$ can be determined using the truncated normal distribution formula:

$$E_t[m_{t+1} \mid \text{distress}] = E_t[m_{t+1} \mid m_{t+1} > \text{threshold}] = \mu_t + \sigma_t \lambda(\alpha_t)$$

where $\alpha_t = (\text{threshold} - \mu_t) / \sigma_t = (T - \mu_t) / \sigma_t$

$$\mu_t = E_t[m_{t+1}]$$

$$\sigma_t = \text{var}_t[m_{t+1}]$$

$$\lambda(\alpha) = \varphi(\alpha) / [1 - \Phi(\alpha)]$$

and $\Phi(\cdot)$ is the standard normal cumulative distribution function. $\lambda(\alpha)$ is called the inverse Mills ratio.

We need to determine a *threshold* T above which banks are considered to be under stress (this threshold should ideally correspond to the states of nature where the CDS payoffs are high). We first assume, somewhat arbitrarily, that the threshold is one historical standard deviation away from the average stochastic discount factor $E_t[m_{t+1}]$ but we show in section III how to extend the model so that the threshold is in line with the estimated probability of default.

C. Estimating the Price of Risk

The remaining variable to estimate is the variance of m_{t+1} , which is linked to the *price of risk* λ_m thanks to the equation:

$$\lambda_m = \frac{Var_t[m_{t+1}]}{E_t[m_{t+1}]}$$

λ_m is called the “price of risk” as the CAPM predicts this factor to be the main driver of excess returns, together with the “quantity of risk”. Indeed, writing equation (4) as

$$1 = E_t[m_{t+1}R_{t+1}] = E_t(m_{t+1})E(R_{t+1}) + \text{cov}(m_{t+1}, R_{t+1})$$

and re-arranging it, the CAPM equation is:

$$E(R_{t+1}) - R_{t+1}^f = \beta_{i,m} \lambda_m$$

where $\beta_{i,m} = -\frac{\text{cov}(m_{t+1}, R_{t+1})}{Var(m_{t+1})}$ is the *quantity of risk* and $\lambda_m = \frac{Var(m_{t+1})}{E(m_{t+1})}$ is the *price of risk*.

In the empirical literature (see Adrian and Moench, 2009 and Adrian and Shin, 2010, for recent applications) the price of risk has been related to the VIX index, or alternatively to the Principal Components (PC) of market yields. The VIX index and the first principal component are strongly correlated (see Coudert and Dex, 2008) so the choice in the literature of one variable or the other does not seem to be crucial.

Method 1. VIX and the Maximum Sharp Ratio

In our second method, we use a property of the price of risk to normalize the VIX (we could do the same with the PC). For a given asset return, the pricing equation implies

$$1 = E(m)E(R^i) + \rho_{m,R^i} Var(R^i)Var(m)$$

$$E(R^i) = R^f - \rho_{m,R^i} \frac{Var(m)}{E(m)} Var(R^i)$$

Since the correlation coefficient ρ_{m,R^i} cannot be greater than 1, we deduce

$$\left| \frac{E(R^i) - R^f}{Var(R^i)} \right| \leq \frac{Var(m)}{E(m)}$$

Theoretically, the price of risk is therefore the maximum Sharpe Ratio attainable for an efficient portfolio.

Historically, Sharpe ratios higher than 3 would be considered very high: we therefore decide to normalize the VIX series (by a factor of 4) in order to scale the VIX index to a series consistent with this property of the price of risk.

Method 2. Principal Components and the Price of Risk

We could also use the PC method proposed by Adrian and Moench (2008) to estimate the price of risk in the U.S. stock markets. Adrian and Moench (2008) applied their method to bond returns and Adrian, Etula and Shin (2010) applied it to currencies, whilst our focus is on S&P stock returns. The method relies on a CAPM equation:

$$R_{i,t+1} - R_{t+1}^f = \beta_i (\hat{\lambda}_{m,t+1} + \tilde{\lambda}_{m,t+1})$$

where $\hat{\lambda}_{m,t+1}$ is the expected component of the market price of risk and $\tilde{\lambda}_{m,t+1}$ is the innovation in the market price of risk.

The expected component of the market price of risk is assumed to be an affine function of the expected PC of stock market returns:

$$\hat{\lambda}_{m,t+1} = \lambda_0 + \lambda_1 \widehat{PC}_{t+1}$$

where λ_0, λ_1 are the coefficients we need to normalize the PC to make convert it into the estimated price of risk.

If one estimates, for each asset the equation,

$$R_{i,t+1} - R_{t+1}^f = a_i + b_i \widehat{PC}_{t+1} + c_i \tilde{PC}_{t+1} + \varepsilon_{t+1}$$

Adrian and Moench (2008) show that $a_i = \beta_i \lambda_0$ and that $\beta_i \lambda_1 = c_i$, two equations that identify the normalization coefficients λ_0 and λ_1 . They apply the model to the Principal Component of bond yields.

III. ENDOGENOUS THRESHOLD

The actual probability of default is deduced from the risk-neutral probability using the

$$\text{formula: } \pi_t = \frac{\hat{\pi}_t}{(1+r_t^f)E_t[m_{t+1} | m_{t+1} > T]} = \frac{\hat{\pi}_t}{(1+r_t^f)(\mu_t + \sigma_t \cdot \lambda(\alpha_t))}$$

where $\lambda(\alpha) = \varphi(\alpha)/[1-\Phi(\alpha)]$ is the inverse Mills ratio and $\alpha_t = \frac{T - \mu_t}{\sigma_t}$.

The threshold T was chosen once and for all, such that $P[m > T] = 0.84$, i.e., the price of risk under distress was one unconditional standard deviation $\bar{\sigma}$ from the unconditional mean $1/(1+\bar{r})$:

$$T = E[m] + \Phi^{-1}[0.84] * \sqrt{\text{var}(m)} = \frac{1}{1+\bar{r}} + \bar{\sigma}$$

where $\Phi^{-1}[0.84] = 1$ and Φ^{-1} is the inverse of the cumulative normal distribution.

The estimation of the market price under risk would gain in coherence if one could choose the threshold T such that the definition of the stress scenario is in line with the probability of distress that is finally estimated. This would also ensure that the threshold is bank-specific and consistent with the probability of default.

Hence, we want to set $T = E[m] + \Phi^{-1}[1-\pi] * \sqrt{\text{var}(m)} = \frac{1}{1+\bar{r}} + \Phi^{-1}[1-\pi]\bar{\sigma}$

where π is the probability of default we are looking for. (Note that a very high threshold means we are looking at unlikely events, i.e., a lower π , which is why $1-\pi$ is the argument in Φ^{-1} .)

The probability of default therefore solves the following non-linear equation:

$$\pi_t = \frac{\hat{\pi}_t}{1 + (1+r_t^f)\sigma_t \cdot \lambda\left(\frac{\frac{1}{1+\bar{r}} - \frac{1}{1+r_t^f} + \Phi^{-1}[1-\pi_t]\bar{\sigma}}{\sigma_t}\right)} \quad (9)$$

or $\pi_t = f(\pi_t)$.

Proposition

There is a unique probability of default that solves equation 9.

Proof

Φ^{-1} and λ are strictly increasing so $f(\pi_t)$ is an increasing function of π_t . Furthermore

$$\lim_{x \rightarrow 0} \Phi^{-1}(x) = -\infty \text{ and } \lim_{x \rightarrow -\infty} \lambda(x) = 0 \text{ implies } f(\pi_t) \xrightarrow{\pi_t \rightarrow 1} \hat{\pi} < 1$$

The limit of $f(\pi_t)$ is 0 since

$$\lim_{x \rightarrow 1} \Phi^{-1}(x) = +\infty \text{ and it is well known that } \lambda(x) \approx x. \text{ This implies } f(\pi_t) \xrightarrow{\pi_t \rightarrow 0} 0$$

We can also show that the derivative of $f(\pi_t)$ is infinite in 0:

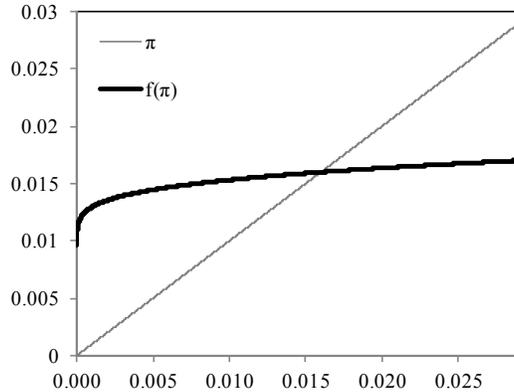
$$\text{Since } \lambda(x) \approx x, \text{ we can rewrite } f(\pi) \approx \frac{1}{a + b\Phi^{-1}[1-\pi]}, \text{ and } f'(\pi) = \frac{b(\Phi^{-1})'[1-\pi]}{(a + b\Phi^{-1}[1-\pi])^2}$$

$$\text{Since } (\Phi^{-1})'[1-\pi] = \frac{1}{\phi(\Phi^{-1}[1-\pi])}, \text{ we find } f'(\pi) = \frac{b}{\phi(\Phi^{-1}[1-\pi])(a + b\Phi^{-1}[1-\pi])^2}$$

Finally, since $\phi(X)(a + bX)^2 \xrightarrow{X \rightarrow +\infty} 0^+$ we proved $f'(\pi_t) \xrightarrow{\pi_t \rightarrow 0} +\infty$.

Therefore, $f(\pi_t)$ is a continuous, strictly increasing, function over $]0, \hat{\pi}]$, with $\hat{\pi} < \pi < 1$ and with a derivative in 0 greater than 1. This shows that the equation $\pi_t = f(\pi_t)$ (equation 8) admits exactly one solution within $]0, 1]$. We plot π_t and $f(\pi_t)$ in Figure 2.

Figure 2. Uniqueness of the Threshold



IV. APPLICATION TO US BANKS DURING THE CRISIS

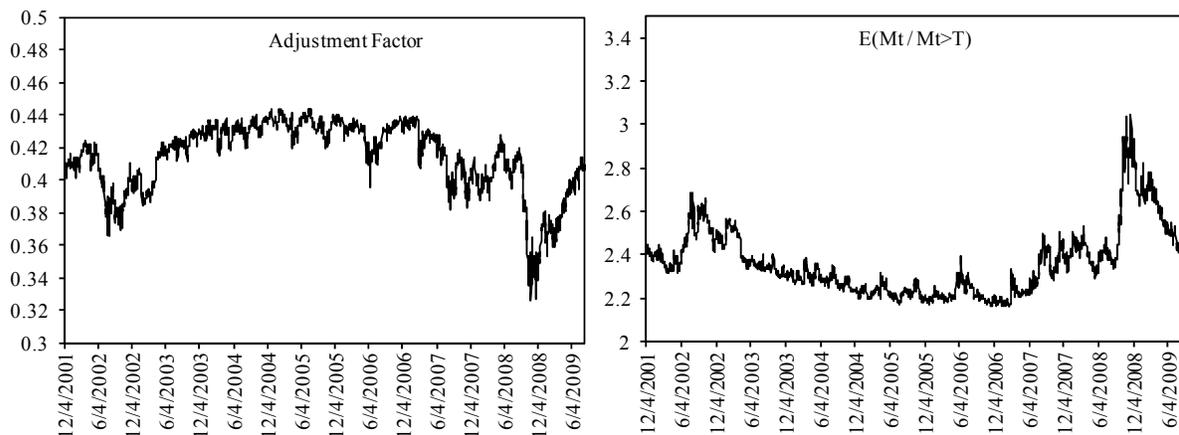
We show in Figure 3 the adjustment factor (between risk-neutral and real probabilities of distress) obtained with the first method described above. In crisis times, the real probabilities of distress tend to be much lower than the risk-neutral probabilities since the “insurance costs” increase as investors seek higher compensation for risk.

We compare our calculated subjective default probability with the Moody’s KMV Estimated Default Frequency. The EDF measure is based on the modeling of equity as a call option on assets (since liabilities have to be paid before shareholders in case of default). The asset value is derived from the market capitalization of the firm, and an estimate of asset volatility is used to model the probability distribution of asset value. When the value of debt, the value assets and the volatility of assets are estimated, the value of equity (which is a call option on the assets) yields the probability of default.

Figure 4 shows the risk-neutral probability of default and the EDF of several US financial institutions, together with our estimate of the default probability under method 2. The risk-neutral probability and associated subjective probability respond to the crisis very quickly while EDF seems to lag. The value, however, of the subjective probability and EDF are more in line than the risk-neutral default probability.

Indeed, we estimated that the market price of risk (for states of nature in which banks would be under stress) increased by 30 percent in the recent turmoil (see Figure 3)

Figure 3. Adjustment Factor and Market Price of Risk



Source: authors’ calculations

Note: the adjustment factor is $\frac{1}{(1+r_f)E_t[m_{t+1} | m_{t+1} > T]}$

V. CONCLUSION

Computing the market price of risk under situations of distress is a crucial element when one needs to use asset prices to assess probabilities of distress. We offered in this paper a simple but theoretically consistent approach to such a calculation. The approach is also free of model assumptions on the shape of the utility function, because it is based solely on an empirical one-factor model. It does not require data on past defaults and therefore is appropriate for estimating probabilities of extreme events—although it still depends on CDS market assessment of risk.

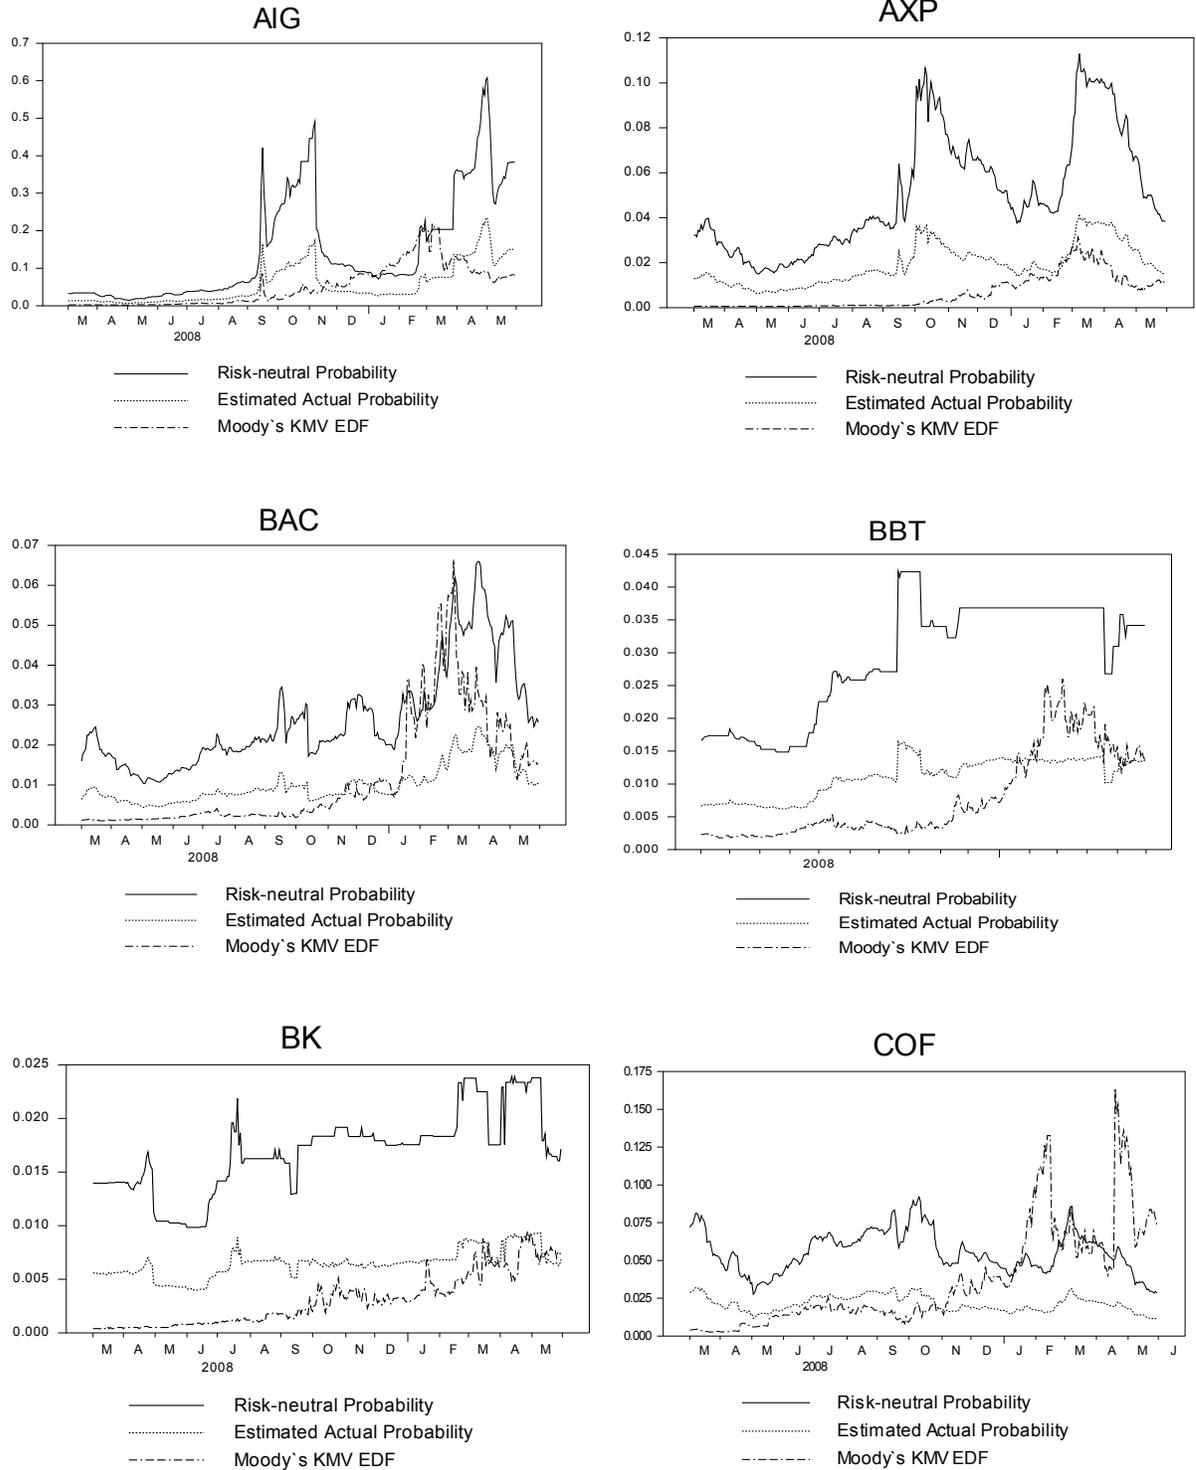
We applied our method to US banks during the subprime crisis, but the method is general enough to be applied to other CDS markets. For instance, using this method, Caceres, Guizzo and Segoviano (2010) analyzed the contribution of risk aversion to the evolution of sovereign spreads in the eurozone.³

The model is based on a one-factor model but a multiple factor model would be better at fitting the data and at capturing the determinant of asset prices. The difficulty lies in matching the risk that we want to measure (the market of price of risk under a situation of distress) with the moments of the pricing factor we can observe. When several factors are assumed to determine pricing, identifying which factor to filter out requires additional understanding of the meaning of the factors and of the meaning of the risk one wants to filter out. In a one factor model, this issue does not arise.

A second assumption in our method is that the price of risk is normally distributed. Relaxing this assumption may also require estimating more factors in the asset pricing model because with a one factor model, only one moment (i.e., the variance) of the price of risk can be identified. Further research is needed to generalize the method to these more complex settings.

³ Caceres, Guizzo and Segoviano (2010) showed that in the aftermath of the subprime crisis, global risk aversion and contagion factors were the main drivers behind increases in euro CDS spreads, but the turning point was in the second half of 2009, when country-specific fundamentals drove CDS in Greece, Ireland and Portugal. Another application of this method for Asia is available in Caceres and Filiz Unsal (2011).

Figure 4. Estimated Probabilities of Default



Source: Moody's KMV, Bloomberg and authors' calculations

REFERENCES

- Adrian, T., and E. Moench, 2008, "Pricing the Term Structure with Linear Regressions," Federal Reserve Bank of New York Staff Report No. 340.
- , E. Etula, and H. Shin, 2010, "Risk Appetite and Exchange Rates," Fed Reserve Bank of New York Staff Report No. 361.
- , and H. Shin, 2010, "Liquidity and Leverage," *Journal of Financial Intermediation*, Vol. 19, pp. 418–437.
- Amato, J., 2005, "Risk Aversion and Risk Premia in the CDS Market," *BIS Quarterly Review*, December, pp. 55–68.
- , and M. Luisi, 2006, "Macro Factors in the Term Structure of Credit Spreads," BIS Working Paper No. 203 (Basel: Bank for International Settlements).
- Berndt, A., D. Rohan, D. Duffie, M. Ferguson, and D. Schranzk, 2005, "Measuring Default Risk Premia from Default Swap Rates and EDFs," BIS Working Paper No. 173 (Basel: Bank for International Settlements).
- Black, F., and M. Scholes, 1973, "The Pricing of Options and Corporate Liabilities," *Journal of Political Economy*, Vol. 7, pp. 637–654.
- Caceres, C., and D. Filiz Unsal, 2011, "Sovereign Spreads and Contagion Risks in Asia," IMF Working paper, forthcoming.
- , V. Guzzo, and M. Segoviano, 2010, "Sovereign Spreads: Global Risk Aversion, Contagion or Fundamentals?" IMF Working Paper 10/120 (Washington: International Monetary Fund).
- Cochrane, J., 2005, *Asset Pricing*, Princeton NJ: Princeton University Press.
- Coudert V., and M. Gex, 2008, "Does Risk Aversion Drive Financial Crises? Testing the Predictive Power of Empirical Indicators," *Journal of Empirical Finance*, Vol. 15, pp. 167–184.
- Jackwerth, J., 2000, "Recovering Risk Aversion from Option Prices and Realized Returns," *Review of Financial Studies*, Vol. 13, pp. 433–451.
- Merton R., 1974, "On the Pricing of Corporate Debt: The Risk Structure of Interest Rates," *Journal of Finance*, Vol. 29, pp. 449–470, 1974.