

On measuring credit risks of derivative instruments

Gregory R. Duffee *

Federal Reserve Board, Mail Stop 91, Washington, DC 20551, USA

Received 15 September 1994; accepted 15 February 1995

Abstract

This paper critically reviews current practices for measuring credit risks of derivative instruments. It argues that there are two major problems with the standard measurement approach. First, it uses models of the stochastic behavior of financial variables while ignoring both their inherent oversimplification and the uncertainty in their parameters. Second, it ignores the correlations among exposures on derivative instruments and the probabilities of counterparty default. This paper demonstrates that these practices can produce large errors in the estimation of distributions of both future credit exposures and future credit losses.

JEL classification: G13; G21; G28

Keywords: Derivatives; Swaps; Credit risk; Risk management

1. Introduction

This paper critically reviews the current state of credit risk measurement of derivative instruments. The proper measurement of credit risks is vitally important to the derivatives market. Market participants need to know how to price credit risk in order to be properly compensated for bearing it. They also need to know how to evaluate the usefulness of mechanisms to reduce credit risk, such as

* Tel.: 202-452-3055; e-mail: gduffee@frb.gov.

transacting with specialized derivative product companies or using collateral. Bank regulators also need accurate measures of the credit risk involved in derivatives so that capital requirements can be set optimally. Regulators may be tempted to compensate for imprecision in the measurement of credit risks by setting very conservative capital standards. However, overly conservative capital standards can unduly restrict the use of derivatives for risk management by setting the shadow price of credit risk too high.

Participants in derivative markets typically have three goals related to credit risk measurement: Estimating future credit exposures, estimating future credit losses, and pricing the risk of default. This paper points out two important flaws in the current measurement methods. The first flaw is that estimates of future exposures, which are based on a Monte Carlo approach, ignore the uncertainties in the models used to generate these estimates.

As an example, I estimate the distribution of future exposure on an interest rate swap. The typical Monte Carlo method using a standard model generates exposures that are roughly 20 percent smaller than the exposures generated by the same model when the uncertainty in the model's parameters is explicitly recognized. The historical distribution of exposures, calculated for the same sample period used to estimate the model's parameters, yields estimates 35 to 60 percent larger than those generated in the standard manner.

The second flaw is that methods used to estimate future losses and price default risk ignore correlations among instruments' exposures and counterparties' probabilities of default. In reality, there are likely to be large correlations among these variables, which implies that joint estimation of exposures and default probabilities is required to estimate credit losses and to price default risk.

For example, if two equally creditworthy counterparties enter into an interest rate swap, the standard method implies that their respective credit risks are roughly equal. However, if both counterparties are typical 'A'-rated corporations, I estimate that the counterparty receiving fixed will have an expected credit loss five times greater than that of the counterparty paying fixed. The difference in expected credit losses is a consequence of the fact that firms are more likely to default in recessions, when interest rates have fallen from earlier levels, and, as a result, fixed-rate receivers are exposed to fixed-rate payers.

Similarly, because correlations among exposures to different counterparties are ignored, the standard method to compute a given upper bound on credit losses associated with a given instrument is essentially useless. I argue that an upper bound on the credit loss associated with a given instrument should be defined as the marginal effect of the instrument on the given upper bound of the aggregate credit loss of the entire portfolio. This marginal effect will depend on the correlation of the instrument's future exposure with the total portfolio's future exposure, and therefore must be estimated within the context of the total portfolio. I present simple examples in which the standard method of estimating upper bounds on credit losses produces estimates that are dramatically wrong.

Section 2 reviews some aspects of the theory of credit risk measurement. Section 3 discusses the current practice of credit risk measurement. Section 4 discusses how to measure future credit losses over some arbitrary period of time shorter than the life of the derivative instrument or instruments. Section 5 concludes.

2. A review of the theory

For the most part, academics and practitioners view credit risk measurement differently. Academics have largely focused on using contingent-claims techniques to price the credit risk associated with derivative instruments, while financial institutions and their regulators have focused on estimating distributions of credit exposures and credit losses. The emphasis on future credit exposures is a consequence of the fact that exposures determine capital requirements, while distributions of future credit losses convey information about the likelihood that credit losses can materially affect a firm's probability of survival. This section discusses the theory underlying these different perspectives of credit risk measurement.

2.1. Pricing default risk

Here I focus on a few papers that are intellectual descendants of the contingent-claims technique pioneered by Black and Scholes (1973). In these papers, complete and arbitrage-free markets are assumed and the value of an uncertain payoff is the expectation, under the equivalent martingale measure, of the payoff discounted with riskfree rates.

In this framework, a model designed to price instruments subject to credit risk must address the following questions. First, what is the process (under the equivalent martingale measure) generating defaults? Second, how are payoffs determined in the case of default? Third, what are the relationships (again, under the equivalent martingale measure) among the variables underlying the default process, the variables underlying the obligated cash flows on the instrument, and riskless interest rates?

The contingent-claims technique was first applied to default risk by Merton (1974), who priced zero-coupon bonds issued by default-risky firms. He assumed that the value of the default-risky firm follows geometric Brownian motion. The payoff to the bond equals the value of the firm if, at maturity, the firm's value is less than the bond's face value. Default cannot occur prior to the maturity of the bond. With this setup, the firm's stockholders own a European option on the value of the firm with a strike price equal to the face value of the bond. He assumed that riskless interest rates are constant, and therefore did not allow the value of the obligated cash flow to vary with the probability of default. This latter assumption

was relaxed by Shimko et al. (1993), who allow for stochastic riskless interest rates using Vasicek's one-factor model of the riskless term structure (Vasicek, 1977), with an arbitrary correlation between the innovations in interest rates and the firm's value.

A practical difficulty with Merton's approach is that for instruments with multiple cash flows, such as coupon bonds, stockholders own compound options on the firm's value (Geske, 1977). Valuing such options generally requires numerical techniques. Cooper and Mello (1991) and Abken (1993) use such techniques to value default-risky interest rate swaps. In these papers, the value of the default-risky firm follows geometric Brownian motion and default occurs at the first payment date at which the firm's value is less than the firm's required payment on the swap. Both papers also allow the processes generating firm values and riskless interest rates to be correlated.

Black and Cox (1976) generalized Merton's model to allow default to occur prior to the maturity date of the bond if the firm's value falls to a given (arbitrary) level. If this level is reached, the bondholders receive the value of the firm. A practical difficulty with both Merton and Black and Cox is that when there are multiple classes of bondholders, the value of the firm must be divided among them when the firm defaults, but Franks and Torous (1994), among others, find that this division does not follow simple priority rules.

Rather than expressly modeling the bargaining process among claimants, Hull and White (1992) and Longstaff and Schwartz (1994) make an assumption that substantially simplifies valuation. Following Black and Cox, they assume that default occurs when the value of the firm falls to a given level, but they also assume that in the event of default, each class of risky debt receives some arbitrarily determined fraction of the riskfree value of the instrument. Their models do not enforce the constraint that the total payoffs to the bondholders equal the remaining value of the firm. Hull and White focus their attention on valuing an instrument assuming independence of the processes generating the firm's value and the riskfree value of the instrument, while Longstaff and Schwartz analytically value an interest rate swap allowing for a nonzero correlation between the processes generating the firm's value and riskfree rates (and therefore the riskfree value of the swap).

In the above models, the value of the firm follows a diffusion process and default occurs when the value of the firm reaches a lower boundary. Roughly speaking, therefore, firms never default unexpectedly. (More precisely, the time of default is a predictable stopping time.) The validity of this implication is questionable. An alternative approach, taken by Madan and Unal (1994), Jarrow et al. (1994), Lando (1994), Duffie and Singleton (1994), and Duffie and Huang (1994) is to use jump processes so that defaults can occur unexpectedly.

Although a detailed summary of the above literature is beyond the scope of this paper, I should note that models which allow dependence between the processes generating firm values and riskless interest rates find that the spread between risky

fixed-rate debt and riskfree fixed-rate debt is strongly positively related to the correlation between innovations in firm value and riskfree rates (Shimko et al., 1993; Longstaff and Schwartz, 1994). When this correlation is positive, the states of the world in which fixed-rate debt is valuable (falling interest rates) are those in which firms are more likely to default. Similarly, the value of an interest-rate swap to the receive-fixed side is inversely related to the correlation between firm values and riskfree rates (Cooper and Mello, 1991; Abken, 1993; Duffie and Huang, 1994). None of these papers make an attempt to parameterize the correlation between the processes generating firm values and riskfree rates with historical data.

2.2. *Estimating distributions of credit exposures and losses*

There are few rigorous discussions of the theoretical issues involved in measuring distributions of credit exposures and losses on derivative instruments in the literature. The following analysis is similar to those in Hull (1989a, Hull (1989b)). For clarity, denote the two counterparties to a given derivative contract as ABC and XYZ. For simplicity I view matters from the perspective of ABC. The cost to ABC of XYZ's default (prior to any recovery in bankruptcy) is the amount required to induce another counterparty to take over XYZ's contractually obligated cash flows. This amount is termed the replacement cost or the credit exposure of the contract. If neither ABC nor the new counterparty have any risk of default, this replacement cost equals the greater of zero and the 'riskless' mark-to-market value of the contract.¹

Much of the literature concerned with credit exposure simply uses the term mark-to-market without the 'riskless' modifier. However, this literature is always referring to the value of the instrument assuming riskless counterparties, not the true mark-to-market value of the instrument, which reflects counterparty credit quality.

To fix notation, denote the time- t 'riskless' mark-to-market value of the derivative instrument (from the perspective of ABC) by p_t . I denote the exposure on the instrument as v_t , which is given by $v_t = \max(0, p_t)$. The distribution of v_t at any time τ can be calculated from the distribution of p_τ , which is in turn determined by the distribution of the underlying financial variables that underlie the derivative instrument.

We need to define a 'credit loss' before we can estimate its stochastic behavior. The literature concerned with estimating distributions of credit losses implicitly

¹ This assumes that default by XYZ does not allow ABC to avoid making payments on derivatives contracts that have a positive mark-to-market value from the perspective of XYZ. If the contracts include provisions for limited two-way payments (walkaway clauses), this assumption may not be valid.

uses the following definition of a credit loss on a derivative instrument. A credit loss occurs when (1) the counterparty defaults and (2) the credit exposure on the instrument is positive. The magnitude of the loss equals the difference between this positive value and the amount recovered from the defaulting counterparty.

It should be clear from this definition that this literature is largely divorced from the contingent-claims valuation literature previously discussed. A reasonable definition of credit loss implied by this literature is the loss incurred when the value of a derivative instrument falls owing to changes in counterparty credit quality. The two definitions are equivalent only when the probability of counterparty default is constant over time. (As we shall see, the assumption of constant default probability is a recurring one in this area.) A discussion of the relative merits of these definitions is deferred until Section 4. The remainder of this section uses the first definition of credit loss.

The risk of credit losses can be formalized by denoting the uncertain fraction of the exposure v_t that is recovered, were default to occur at time t , as δ_t . The credit loss on this instrument is described by

$$\text{credit loss} = \begin{cases} 0, & \text{if XYZ does not default;} \\ (1 - \delta_\tau)v_\tau, & \text{if XYZ defaults at time } \tau. \end{cases} \quad (1)$$

A similar formulation holds for a portfolio of contracts between ABC and XYZ. Consider a portfolio with mark-to-market values $p_t^1, p_t^2, \dots, p_t^n$ and corresponding exposures $v_t^1, v_t^2, \dots, v_t^n$. The distribution of credit losses depends on whether the contracts are covered by a legally enforceable netting agreement. If netting is effective, credit losses on the portfolio are described by

$$\text{credit loss} = \begin{cases} 0, & \text{if XYZ does not default;} \\ (1 - \delta_\tau) \max \left[0, \sum_{i=1}^n p_\tau^i \right], & \text{if XYZ defaults at time } \tau. \end{cases} \quad (2a)$$

If netting is not allowed, credit losses on the portfolio are described by

$$\text{credit loss} = \begin{cases} 0, & \text{if XYZ does not default;} \\ (1 - \delta_\tau) \sum_{i=1}^n v_\tau^i, & \text{if XYZ defaults at time } \tau. \end{cases} \quad (2b)$$

An issue that is generally ignored in this literature (as opposed to the contingent-claims literature) is the proper way to discount future credit losses. Such discounting should be tailored to the problem at hand. If our objective were to determine the current market value of future credit losses, we would proceed as in the contingent-claims literature. Alternatively, we might be interested in the probability that credit losses do not exceed ABC's assets. If so, we should discount credit losses at the (stochastic) rate of return to those assets. Throughout

this paper I work with undiscounted losses, to avoid the complications introduced by discounting.

Implicit in Eqs. (1), (2a), and (2b) is the idea that credit loss distribution depends on the joint distribution of the exposure, the probability of default of XYZ, and the fraction of losses recovered in bankruptcy. This idea is made more explicit in the expression for the expected credit loss. This expected loss, conditional on information available at time t , is given by the probability of default at some point during the life of the contract multiplied by the expected unrecovered exposure at the time of default:

$$E_t(\text{credit loss}) = \text{prob}_t(\text{XYZ defaults}) \cdot E_t[(1 - \delta)v | \text{XYZ defaults}] \quad (3)$$

In the case of a portfolio of contracts this expectation can be written as (4a), if netting is allowed, or (4b), if netting is not allowed:

$$E_t(\text{credit loss}) = \text{prob}_t(\text{XYZ defaults}) \cdot E_t\left((1 - \delta) \max\left[0, \sum_{i=1}^n p^i\right] | \text{XYZ defaults}\right) \quad (4a)$$

$$E_t(\text{credit loss}) = \text{prob}_t(\text{XYZ defaults}) \cdot E_t\left((1 - \delta) \sum_{i=1}^n v^i | \text{XYZ defaults}\right) \quad (4b)$$

The expected exposures (and recovery ratios) on the right-hand sides of Eqs. (3), (4a), and (4b) are conditional on both time- t information and the default of XYZ at some time during the life of the instrument. If the credit exposure on the instrument is correlated with the probability of XYZ's default, this expectation generally will differ from both the unconditional expected exposure and the expected exposure conditioned on time- t information.

Few papers attempt to empirically estimate credit losses on derivative instruments. Hull (1989b) calculates expected credit losses on interest-rate swaps and currency swaps relative to expected credit losses on loans under the assumption that the riskfree term structures are flat and instantaneous changes in rates are log-normally distributed. A key parameter in his calculations is the assumed correlation between riskfree rates and the default risk of firms. The credit loss to the receive-fixed side in an interest-rate swap is positively related to this correlation: As the correlation increases, the states of the world in which fixed-rate debt is valuable are those in which firms are more likely to default.

Belton (1987) considers the distribution of credit losses on a portfolio of interest-rate swaps with many counterparties, where the distribution of movements in the yield curve are generated by a vector autoregressive model. Although he recognizes that the distribution of losses depends on the correlation between interest rates and default probabilities, he assumes a zero correlation in his

simulations. As we see in the next section, his assumption is shared by practitioners involved in measuring and managing credit risks on derivative instruments.

3. A review of the practice

A review of the current practice of credit risk measurement is confounded by three facts. First, each participant in the derivatives market has its own procedures for measuring credit risks, with varying degrees of sophistication. Second, 'current' practice is changing fairly quickly over time, as the general level of sophistication of participants grows. Third, the most sophisticated techniques developed by major dealers are generally proprietary. 'Current practice', as described in this paper, refers to the techniques used by the more sophisticated participants, based on published research, press reports, industry seminars, and discussions with various industry professionals involved in designing techniques for measuring credit risk.² Moreover, the Group of Thirty (1993) has recommended that all dealers and end-users adopt the 'current practices' as described here.

This section is divided into several parts. A brief overview of current practice is found in Section 3.1. Section 3.2 provides a more detailed description of Monte Carlo methods for estimating distributions of credit exposures; Section 3.3 criticizes this Monte Carlo technique as currently practiced. Section 3.4 describes the standard methods used to compute estimated credit losses and an upper bound on these credit losses. Sections 3.5 and 3.6 criticize these standard methods for measuring expected credit losses and an upper bound on credit losses, respectively.

3.1. An overview

Measures of future credit exposures are estimated as follows. Distributions of exposure at many discrete points (say, τ_i , $i = 1, \dots, K$) over the life of the instrument or portfolio are computed. These distributions are sometimes conditioned on the current state of financial variables, but not conditioned on the probability of counterparty default. For each of the K distributions, a particular summary statistic (say, mean or some upper bound) is then computed. These K summary statistics are then averaged to produce a single measure of future exposure (an average mean or an average upper bound) for the instrument or portfolio.

Measures of future credit losses are produced by multiplying measures of future

² Good descriptions of the approaches taken by rating agencies in measuring derivatives' credit risk are Standard & Poor's (1992) and Gluck and Clarkson (1993).

credit exposure by the cumulative probability of counterparty default over the life of the instrument. This result is further multiplied by the fraction of loss expected to be unrecovered in bankruptcy. For example, if the measure of future credit exposure is the 95th percentile upper bound, this multiplication produces the 95th percentile upper bound on credit losses.

Default risk is often priced using measures of expected credit losses instead of using contingent-claims techniques. Frequently, the price effect is simply the difference in the counterparties' expected credit losses as estimated above, plus or minus an arbitrary spread. More sophisticated practitioners use option pricing methods in which the underlying financial variables and the probability of counterparty default evolve independently.

These approaches do not incorporate any possible correlation between future exposure and the probability of counterparty default. However, ad hoc corrections to this approach are sometimes made. The most common correction is made for interest-rate swaps, as described in Sorensen and Bollier (1994). When the term structure is upward sloped, the first net payment is made to the receive-fixed side. After this payment the contract is expected to be in-the-money to the pay-fixed side. Because defaults tend to occur later rather than earlier, the pay-fixed side is perceived to face more credit risk than is the receive-fixed side. However, as we shall see, this 'correction' is inappropriate; the receive-fixed side faces much more credit risk than the pay-fixed side because defaults tend to occur following declines in interest rates. We now examine these measurement practices in more detail.

3.2. The standard method of estimating future exposures

Distributions of future exposure are typically generated by the technique of Monte Carlo simulation. The standard approach is to choose a particular parameterization of the financial variables of interest. For example, to calculate the distributions of future exposures for an interest-rate swap, I might choose a one-factor Cox, Ingersoll, and Ross (CIR) model of the riskfree term structure (Cox et al., 1985). The parameters used in the simulation are usually calibrated with historical data. For example, the CIR model might be estimated on historical observations of the U.S. dollar term structure.

The resulting model is used to generate randomly a hypothetical time path of the relevant financial variables. The replacement values of an instrument or portfolio are calculated at each point along each time path. This exercise is repeated thousands of times. The resulting distributions of replacement costs are used to calculate expectations and confidence bounds for exposures.

A large number of researchers have generated statistics concerning the distribution of exposures for various derivative instruments. The general focus has been on interest-rate swaps and currency swaps. A partial list includes Arak et al. (1986), Muffet (1987), Bank of England and Federal Reserve Board Staff (1987),

Ferron and Handjinicolaou (1987), Whittaker (1987), Bankers Trust (1987), Neal and Simmons (1988), Duffee (1992), Hendricks and Barker (1992), Gilberti et al. (1993), and Bond et al. (1994). In addition, most major banks have staff engaged in estimating these distributions.

I illustrate this Monte Carlo approach in the context of interest-rate swaps with a one-factor CIR model of nominal yields. In the CIR model, the evolution of the instantaneous riskless interest rate follows the diffusion process $dr = \kappa(\alpha - r)dt + \sigma\sqrt{r} dZ$. I estimate the following discrete-time version:

$$\begin{aligned} \Delta r_{t+1} &= \kappa(\alpha - r_t) \Delta t + e_t, & e_{t+1} &= \sigma\sqrt{r_t} \Delta t \eta_{t+1}, & E(\eta_{t+1}) &= 0, \\ E(\eta_{t+1}^2) &= 1. \end{aligned} \quad (5)$$

Eq. (5) is estimated with the generalized method of moments (GMM), using monthly observations of the three-month T-bill yield from the Center for Research in Security Prices (CRSP) as a proxy for the instantaneous interest rate. The estimation period is March 1959 through December 1992. The moment conditions are $\bar{e}_t = e_{t+1} e_t = e_t^2 - \sigma^2 r_t = 0$. The results are below, with standard errors in parentheses:

$$\Delta r_{t+1} = \underset{(0.223)}{0.268} \left(\underset{(0.013)}{0.063} - r_t \right) \Delta t + \underset{(0.006)}{0.082} \sqrt{r_t} \Delta t \eta_{t+1}$$

The CIR model is closed with a market price of interest-rate risk, which I assume is zero. This assumption is in accord with the results in Chen and Scott (1993). They estimate the one-factor CIR model using maximum likelihood and cannot reject the hypothesis that the market price of risk is zero. Zero-coupon bond prices can then be calculated as functions of the current interest rate, as described in Cox et al. (1985).

The CIR model has been estimated with much more sophisticated techniques such as in Chen and Scott (1993). However, for the purpose of estimating distributions of future exposure, the approach taken in this paper is more rigorous than most. I return to the issue of econometric sophistication below.

Here I restrict my attention to a five-year U.S. dollar plain vanilla interest-rate swap with semiannual payments. I first estimate distributions of exposure, from the perspective of the receive-fixed side, for a new swap. I estimate both unconditional and conditional distributions, where the conditioning information is the shape of the yield curve at the time the swap is initiated. In a CIR model, the shape of the yield curve at time t is summarized by the model parameters and the time- t instantaneous interest rate. The first panel in Fig. 1 displays expected exposures for distributions that are (1) unconditional (the solid line), (2) conditioned on the initial interest rate equal to 12% (the dot-dashed line); and (3) conditioned on the initial interest rate equal to 3.5% (the dotted line). The unconditional distributions are estimated by first generating a very long time series of interest rates using the parameterized model, then randomly selecting values




Fig. 1. Expected and 95th percentile upper bounds on credit exposures (in % of notional principal) for the receive-fixed side of a five-year U.S. dollar interest-rate swap. The horizontal axis measures years since the initiation of the swap. Distributions are calculated with Monte Carlo simulation of a Cox–Ingersoll–Ross model. The solid lines use the unconditional distribution of exposures. The dotted lines assume the initial interest rate equals 12% and the dot-dashed lines assume the initial interest rate is 3.5%.

from this time series to use as initial interest rates. The second panel displays the 95th percentile upper bound for these distributions.

The information in distributions of future exposures is compressed into a few summary statistics. These statistics are defined here in the context of the distribution of exposures for the interest-rate swap examined in Fig. 1.

Average expected exposure is defined as the mean (over time) of expected exposure. For example, the unweighted mean of the solid line in the first panel of Fig. 1 is the unconditional average expected exposure, while the unweighted means of the other lines are average expected exposures conditioned on the relevant initial interest rate. *Average maximum exposure* is defined as the mean (over time) of some upper bound on exposure. ‘Maximum’ is a misnomer because exposure can obviously exceed this maximum. If we choose an upper bound of 95%, then the average maximum exposures for the distributions in the second panel of Fig. 1 are the unweighted average values of the lines in this figure. *Peak expected exposure* and *peak maximum exposure* are defined as the respective peaks of the lines in the first and second panel in Fig. 1.

Fig. 1 demonstrates that incorporating the initial shape of the term structure into calculations of future exposure can significantly affect the perceived risk of a new interest-rate swap. When the initial interest rate is 12%, the yield curve is downward sloped and interest rates are expected to fall. Therefore the floating rate

is initially higher than the fixed rate and the initial net payments are made by the receive-fixed side. After these initial payments, the swap is expected to be in-the-money to the receive-fixed side. Therefore expected future exposures (from the receive-fixed perspective) conditioned on an initial interest rate of 12% are roughly 2.3 times larger than unconditional expected future exposures. Conversely, expected future exposures conditioned on an initial interest rate of 3.5% are roughly 0.6 times as large as unconditional expected exposures. When estimating distributions of future credit exposures of seasoned swaps, it is even more important to condition on the history of interest rates because the swaps typically will have moved into or out of the money.

3.3. Criticism of the standard method of estimating future exposures

Future values of credit exposure on a given derivative instrument are unknown because the future values of the underlying variables that determine the cash flows are unknown. The above-described Monte Carlo approach implicitly assumes that the only source of uncertainty concerning these underlying variables is that we do not know the realization of the path generated by the model. In reality, however, there is another important source of uncertainty: The true model generating these paths is unknown.

We would like to have some sense of the importance of this uncertainty. If the model used in the Monte Carlo estimation is assumed to be true but the parameters are recognized to be unknown, we can estimate credit exposures with a different type of Monte Carlo simulation. To generate a single time path of underlying variables, we first randomly choose a set of parameters, then use the parameters (in conjunction with the model) to generate the time path. This technique requires a joint distribution from which to draw the parameters. One (albeit imperfect) way to construct such a distribution is to use standard econometric techniques to estimate the model. The distribution is then assumed to be multivariate normal, with a mean equal to the parameter estimates and variance–covariance matrix equal to the estimated variance–covariance matrix from the econometric estimation. (Of course, this procedure implies that the parameter distribution is known, when in fact it is estimated.)

To illustrate the differences between the standard Monte Carlo approach and this random-parameters Monte Carlo approach, I recalculated unconditional distributions of exposure on a five-year interest-rate swap. The same CIR model is used, but parameters are chosen from a multivariate normal distribution with a mean equal to the estimated parameters and a variance–covariance matrix given by the GMM estimation. (This approach is less accurate than it might be. The variance of interest rates must be nonnegative, hence the distribution from which σ is drawn should not be normal.)

A comparison of these two approaches is displayed in Fig. 2. The first panel reports expected exposure, while the second panel reports the 95th percentile

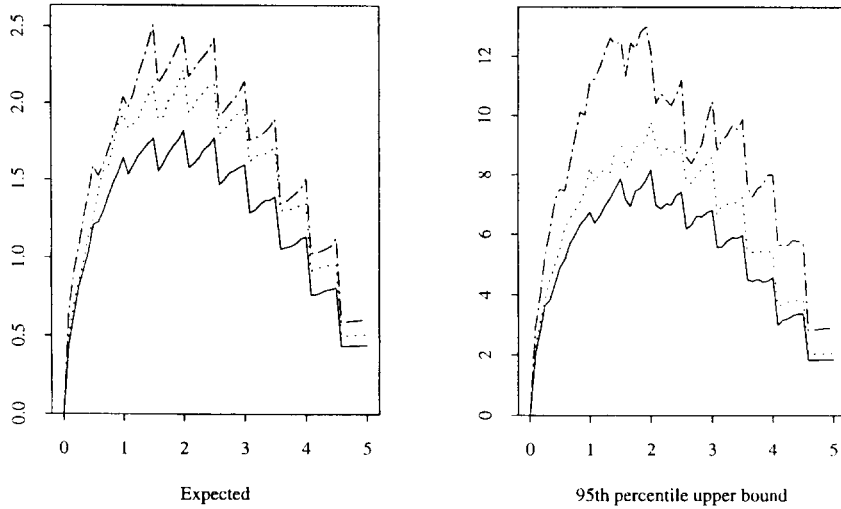


Fig. 2. Expected and 95th percentile upper bounds on unconditional credit exposures (in % of notional principal) for the receive-fixed side of a five-year U.S. dollar interest-rate swap. The horizontal axis measures years since the initiation of the swap. The solid lines use Monte Carlo simulation of a Cox–Ingersoll–Ross model with fixed parameters. The dotted lines use Monte Carlo simulation of the same model with random parameters. The dot-dashed lines use bootstrapping.

bound on exposure. In both panels, the solid lines are the distributions given by the standard approach — they are duplicates of the solid lines in Fig. 1. The dotted lines are distributions given by this alternative approach. (Ignore the dot-dashed line for the moment.) The dotted lines are approximately 19% higher than the solid lines, indicating that the standard Monte Carlo approach underestimates credit exposures because it assumes we know more than we really do. Moreover, differences between the standard approach and the random-parameters approach are likely to be far more pronounced in typically used models than in the one-factor CIR model used here. In practice, term-structure models such as Ho and Lee (1986) use a large number of parameters in order to fit precisely the current term structure. These parameters are highly unstable over time.

This alternative Monte Carlo approach likely still underestimates both mean credit exposures and upper bounds on credit exposures because it ignores the fact that we do not know the true model. To properly account for this uncertainty, we should randomize over possible models. Unfortunately, we cannot even define the universe of possible models, much less determine a probability measure over this universe. Therefore I take another approach.

I use bootstrapping methods instead of Monte Carlo simulations to estimate distributions of future credit exposure. This method requires that there be sufficient historical data to allow us to choose a large number of different actual paths. Therefore it will not be useful in generating exposure distributions for derivative

instruments that are priced off of variables for which there are no long time series. However, there are long time series of observations on the U.S. dollar yield curve. I therefore use the bootstrap method to estimate the unconditional distribution of exposure on a new five-year U.S. dollar interest-rate swap.

I use CRSP data to construct month-end term structures for February 1959 through December 1992. For each of these 407 months, I construct sixty different yield-to-maturities corresponding to zero-coupon bonds that mature in one month, two months, and so on up to sixty months. I then use these data to calculate, for each month, the fair-value fixed rate for a newly issued five-year swap with semiannual payments. For each of the 348 months between January 1964 and December 1992, I next calculate the mark-to-market value, and therefore the credit exposure, of a five-year interest-rate swap with one or more months to maturity. (A five-year swap that initiated at the end of February 1959 would have matured in February 1964.)

The dot-dashed lines in Fig. 2 display the historical distributions of future exposure on the swap. The distributions have much higher means and confidence bounds than do the distributions calculated with the Monte Carlo techniques. The average expected exposure using the historical approach is over 33% larger than the corresponding average using the standard Monte Carlo approach, while the average 95th percentile exposure using the historical approach is almost 60% larger than the corresponding average for the standard Monte Carlo approach.

The primary reasons for the larger exposure using the bootstrap method is that the actual distribution of changes in interest rates exhibits both fat tails and short-term persistence in volatility. These features are poorly captured by the CIR model. This is illustrative of a general problem. Every model is too simple – i.e., it has omitted variables. Two common oversimplifications are to ignore time-varying volatility and to ignore imperfect correlations among various financial variables (assuming that either the correlation is perfect or that there is no correlation at all).

Unfortunately, although bootstrapping is a very useful check on the accuracy of a given model, it is not a substitute for a model-based approach. The problem is that, for risk management purposes, distributions of future exposure should be conditioned on the current state of the relevant financial variables. As we saw in Section 3.2, conditional distributions are often much different from the unconditional distribution. However, there is not enough historical data to generate distributions of exposures conditioned on particular values of the relevant financial variables without using a model.

For example, consider estimating distributions of future exposure on a new interest-rate swap given the current term structure. With a reasonably sophisticated model, the current term structure can be used as an input to the model. With the historical approach, empirical distributions of future exposure must use only those realized term-structure paths that have an initial term structure similar to the current term structure. Even with decades of data, this set of paths is likely to be

far too small to produce sufficiently precise estimates of distributions of future exposures. Only model-based Monte Carlo estimation can generate the necessarily large number of paths.

A final problem shared by all of these approaches is that they assume that the future will be like the past. More precisely, accuracy in model calibration requires that the future data be generated by the same process that generated historical data. Given rapid changes in financial markets over time, this equivalence may not hold. For example, imagine that a ten-year oil swap had been initiated in 1969. At that time, any model's 99th percentile confidence bound on future exposures would have grossly underestimated actual exposures.

Perhaps the best defense against errors induced by such regime changes is a familiarity with the frequency of regime changes in a wide range of financial variables. Although oil price data from 1950 through 1970 could not have predicted the great volatility of the 1970s, commodity prices often exhibit periods of stability followed by periods of great volatility. These periods of volatility are often associated with financial panics or international crises (stock prices in 1929, grain prices at the beginning of World War I). Recognizing the possibility of such periods of high volatility might substantially improve estimates of the tails of exposure distributions.

3.4. The standard method of estimating credit losses and pricing default risk

Armed with measures of future credit exposure, credit losses associated with a contract or portfolio of contracts with a single counterparty typically are estimated as follows. The expected credit loss is the product of the average expected exposure and the estimated cumulative probability of counterparty default (over the life of the contract or portfolio of contracts). The 'maximum' credit loss is usually defined as the product of the average maximum exposure and the estimated cumulative probability of counterparty default, although sometimes peak maximum exposure is used instead of average maximum exposure.

I do not explicitly discuss the option-like methods to price default risk because in practice, they are very similar to methods used to measure expected credit losses. Although in principle, expected credit losses depend on the true probability distributions of the underlying variables while option-like methods require the pseudo-risk-neutral distributions, in practice risk neutrality is widely assumed when calculating expected credit losses.

There are a variety of ways to estimate the probability of counterparty default. The rating agencies are standard sources for average default probabilities and recovery rates for publicly rated corporations, in addition to individual firm credit ratings. Moody's data is summarized in Fons et al. (1994), while S&P's data is summarized in Brand et al. (1994). Techniques used to forecast default probabilities for individual firms can be found in Altman et al. (1977). Proprietary techniques that extract information about underlying asset values and volatility

Table 1
Cumulative default probabilities by initial Moody's credit rating, 1970–1993

Initial rating	Probability of default by end of year:									
	1	2	3	4	5	6	7	8	9	10
Aaa	0.000	0.000	0.000	0.000	0.001	0.002	0.003	0.004	0.006	0.007
Aa	0.000	0.000	0.001	0.002	0.003	0.004	0.005	0.006	0.008	0.009
A	0.000	0.001	0.003	0.005	0.006	0.008	0.011	0.013	0.016	0.020
Baa	0.002	0.005	0.009	0.015	0.020	0.025	0.031	0.037	0.044	0.050
Ba	0.018	0.044	0.069	0.094	0.118	0.138	0.153	0.167	0.181	0.195
B	0.083	0.148	0.204	0.248	0.284	0.319	0.343	0.367	0.384	0.400

Source: Fons et al. (1994).

from stock prices are commercially marketed. The market price of default risk, which incorporates the probability of default, uncertain recovery rates, and the market prices of these risks, can be extracted from the premium of publicly traded risky debt over riskless debt, although Altman's evidence (Altman, 1989) that excess yields on corporate bonds are consistently positive suggests that market prices be used with caution. In addition, bank credit departments produce internal estimates of counterparty credit risk.

Below I make two criticisms of this approach to measuring credit losses. First, the calculations of expected credit losses ignore correlations between default probabilities and exposures. Second, statistics such as an upper bound on credit losses associated with a single instrument or portfolio of instruments cannot be calculated usefully without reference to the total portfolio of the institution.

3.5. Criticism of the standard method of measuring expected credit losses

There are three reasons why exposures and default probabilities generally will be correlated. First, for a given firm, the probability of default over a given time interval Δt will vary over the life of a derivative instrument. The cumulative default probabilities for various Moody's rating categories over 1970–1993 are reported in Table 1. The table documents that the probability of an 'A'-rated counterparty defaulting within two years is 0.1%, while the probability of default during the subsequent two years is 0.4%. As can be seen in the table, this pattern is typical for originally investment-grade firms (but not for originally speculative-grade firms).

Therefore, for a typical investment-grade counterparty, the risk of a credit loss on a derivative instrument depends more on exposures near the end of the instrument's tenor than on exposures near the beginning. Consider, for example, two different contracts with equal average expected exposure over their lives. If one contract has much of its exposure concentrated in the early portion of its life (such as an interest-rate swap), and the other contract has much of its exposure

concentrated in the late portion of its life (such as a currency swap), the latter contract will have a greater risk of credit loss.

The second reason why exposures and default probabilities will be correlated is that default probabilities are correlated with the business cycle (Fons et al., 1994), which is tied to the macroeconomic variables that underlie most derivative instruments. For example, Estrella and Hardouvelis (1991) find that the slope of the term structure is a predictor of future economic activity.

The third reason is that firms generally enter into derivative contracts that have mark-to-market values which are correlated with firm- or industry-specific shocks to firms' survival probabilities. For example, end-users typically enter into contracts to either hedge other positions or to speculate in areas in which they have special knowledge. The situation with major derivative dealers is somewhat different. Dealers transact with many different types of end-users, and therefore are unlikely to have derivative books that are extremely sensitive to the movements in a single financial variable, such as the price of oil. However, the default of a major dealer is more likely to occur during a time of larger-than-average volatility across many financial markets, in which the dealer has made a series of bad bets. It is precisely at this time that even well-diversified derivatives portfolios are likely to be characterized by large exposures.

These links between exposures and default probabilities have been recognized by both academics and market participants.³ However, there are large practical problems involved in estimating these links. For any particular counterparty and derivative instrument, estimating the correlation between default and exposure would be very time-consuming, very difficult, and very imprecise. One simplification is to use aggregate default behavior as a proxy for individual default behavior. This simplification is, of course, inappropriate if the counterparty in question does not look like the 'typical' counterparty. In addition, it cannot incorporate the effect of the counterparty's derivatives activities on the probability of counterparty default – this effect depends on the level and volatility of the counterparty's assets.

If we decide to use aggregate default data as a proxy for individual default experience, it would be fairly simple to modify current practice to incorporate the dependence of aggregate default probabilities on time. The simplicity owes to the fact that time is nonstochastic, so no additional modeling needs to be done.

For example, recall that when the 'expected' credit loss was calculated in Section 3.4, an unweighted average of the expected exposures at each point in time was multiplied by the cumulative probability of default. An 'expected' credit loss that incorporates the dependence of default probabilities on time is calculated almost identically. A *weighted* average of expected exposures at each point in time is computed instead of an unweighted average. The probabilities of default over

³ See, e.g., Belton (1987), Hull (1989a), Gluck and Clarkson (1993), and Lucas (1994).

each interval in time are used as weights. The resulting weighted average is then multiplied by the cumulative default probability and the expected fraction of losses unrecovered in bankruptcy to determine the 'expected' credit loss. Alternatively, an ad hoc approach can be taken, as described in Sorensen and Bollier (1994).

Incorporating correlations between default probabilities and stochastic financial variables is more difficult because it requires knowledge of the joint stochastic behavior of defaults and the relevant financial variables. Unfortunately, there are no empirically calibrated models of the joint behavior of typical financial variables, such as interest rates or equity prices, and the probability of counterparty default. If we had such a model, we could use Monte Carlo simulation to generate a large number of time paths of exposures. At every point along each path there is some probability of counterparty default. Expected credit losses (prior to any recovery in bankruptcy) are calculated by multiplying the exposures at each point by the probability of counterparty default at that point. The model can be made more realistic by including an uncertain fraction of replacement cost recovered in bankruptcy proceedings.

I do not attempt to construct such a model here. Instead, I use a bootstrapping method that relies on the historical behavior of interest rates and corporate defaults. This simulation illustrates the idea that ignoring the joint behavior of defaults and exposures can lead to substantial mismeasurement of credit risks. However, the data analyzed here are too limited for the results to be accepted uncritically.

For this exercise I use the aggregate default experience of corporations rated by Moody's over the period 1971–1992. Consider the cohort of firms rated 'A' by Moody's in January of year t . In Fons et al. (1994), Moody's reports the fraction of these firms that defaulted at the end of year t , $t + 1$, and so on. Moody's reports yearly aggregates, while this exercise uses monthly data. I therefore interpolate the Moody's data in two ways. First, I assume that those firms that default during year $t + i$, $i \geq 0$, do so uniformly within year $t + i$. Second, because Moody's does not report results for cohorts formed in months other than January, I linearly interpolate the Moody's numbers to estimate the default experience for cohorts formed in these months.

I examine the distribution of credit losses on two types of contracts: a five-year U.S. dollar interest-rate swap and a five-year U.S. dollar/German mark currency swap. The credit exposures on these contracts are calculated using actual U.S. and German risk-free term structures, as well as the actual spot exchange rate, over the period 1971–1992. The counterparties are assumed to all be 'A'-rated at the time the swaps were initiated. For simplicity, I do not incorporate recovery in bankruptcy in this simulation. However, it could be added without much additional complexity.

With the historical data used here, there are 204 different five-year time paths of the financial variables of interest. The first path is the realized path over January 1971–January 1975; the last path is the realized path over December

Table 2

A comparison of different methods of calculating expected credit losses on certain derivative instruments

Contract	Not conditioned on default behavior		Conditioned on counterparty default	
	Expected exposure	Implied expected credit loss	Expected exposure	Implied expected credit loss
Five-year US\$ interest-rate swap				
Receive-fixed	2.29	0.013	3.79	0.021
Pay-fixed	2.13	0.012	0.66	0.004
Five-year US\$/DM currency swap				
Pay-\$	18.22	0.103	31.25	0.176
Receive-\$	5.23	0.029	3.65	0.021

Distributions of credit exposures are calculated with actual U.S. and German risk-free term structures and the spot \$/DM exchange rate, over 1971–1992. Default behavior over the same period is for firms originally rated 'A' by Moody's. Expected losses are calculated before any recovery in bankruptcy. All figures are expressed as a percent of notional principal.

1987–December 1992. Corresponding to each of these paths is the actual or estimated default experience of Moody's-rated corporations.

Table 2 reports expected exposures and corresponding expected credit losses for both sides of the interest-rate swap and the currency swap. These expected exposures are first calculated independently of the probability of counterparty default, then conditional on counterparty default.

If the correlation between interest rates and default probabilities is ignored, the expected credit losses on receive-fixed and pay-fixed sides of an interest-rate swap appear similar. The expected credit loss to the receive-fixed side is 1.3 basis points of notional principal, while the expected credit loss to the pay-fixed side is 1.2 basis points. However, when the joint distribution of interest rates and default probabilities is used to estimate credit losses, the credit risk of the receive-fixed side is much larger than the credit risk of the pay-fixed side.

Over the past 20 years, the probability of default by an initially 'A'-rated counterparty was negatively correlated with changes in interest rates – defaults rose and interest rates fell as the economy entered recessionary periods. The value of receiving fixed is high when floating rates fall, which is precisely when counterparties are more likely to default. This is reflected in the expected credit losses for the receive-fixed and pay-fixed sides reported in Table 2. The expected credit loss for the receive-fixed side is approximately two basis points of notional principal, over 65% larger than the expected loss calculated independently of the probability of counterparty default. It is also over five times larger than the expected credit loss (based on the joint probability) for the pay-fixed side.

For currency swaps, even if the correlation between the exposure and counterparty default probability is ignored, the expected credit loss to the pay-dollar side (10 basis points of notional principal) is much larger than the expected credit loss to the receive-dollar side (3 basis points). This disparity reflects the long unexpected depreciation of the dollar over this period. The receive-dollar side was seldom in-the-money. This disparity also points out one of the difficulties of the bootstrapping method to estimating credit risks. In order to have substantial confidence in the results, the data must cover all of the regimes one is likely to encounter in the future. It is unlikely that the behavior of the dollar/DM exchange rate over past 20 years will characterize the future behavior of this exchange rate.

For the currency swap as well as for the interest-rate swap, the expected exposures calculated conditional on counterparty default are substantially different from the expected exposures calculated independently of counterparty default. The default-conditioned exposure to the pay-dollar side is over 70% larger than the default-independent exposure, while the default-conditioned exposure to the receive-dollar side is over 30% smaller than the default-independent exposure. These differences owe to the fact that over this 20 year period, default probabilities were higher when the dollar was falling.

These results should be regarded as illustrative rather than precise measurements of the risk of credit losses. The typical corporation rated by Moody's may not look like the typical counterparty in a bank's derivatives book. On a gross basis, the largest exposures are to other banks and financial institutions instead of corporate end-users. The default behavior of these institutions (including foreign financial institutions) is likely to be different from the default behavior of the average nonfinancial institution. (However, calculation of exposures on a net basis typically reduces the relative importance of financial institution exposures in a typical bank's derivatives portfolio.) More work needs to be done to characterize the relationship between relevant financial variables, such as interest rates and equity prices, and the default behaviors of financial and nonfinancial institutions.

3.6. Criticism of the standard method of measuring upper bounds on credit losses

Here I make two points concerning confidence bounds on credit losses. First, an upper bound on credit losses is a meaningful concept only within the context of a portfolio of instruments with *various* counterparties – it makes little sense to discuss an upper bound for a single counterparty in isolation. Second, the typical approach used to calculate upper bounds on exposure to a given counterparty is useful only under very strict assumptions about the structure of this portfolio that cannot conceivably be met in practice.

An upper bound on credit losses is not a useful measure in the context of a single counterparty because the probability of a single counterparty defaulting is very small. For example, consider a five-year contract (or portfolio of contracts with a maximum life of five years) with a single typical 'A'-rated counterparty.

Table 1 reports that there is less than a 1% chance that the counterparty will default during the life of the contract. Hence both the 95th and 99th percentile upper bounds on credit losses on the contract are precisely zero.

Of more interest is the marginal effect of a particular derivative instrument on an upper bound on aggregate credit losses to a portfolio of instruments with various counterparties. Under certain restrictive assumptions, this marginal effect is independent of the composition of the portfolio and is a multiple of the average expected exposure on the instrument. I am unaware of any arguments justifying measuring this marginal effect by the typical method of multiplying the average upper bound exposure on the instrument by the probability of counterparty default. This approach is probably only appropriate if the distributions of instruments' exposures are such that their average upper bounds are proportional to their average expectations.

Hull (1989a) lists assumptions that are sufficient to justify the use of a multiple of average expected exposure. He notes that they are 'questionable' in practice. First, the portfolio must be very large, so that the Central Limit Theorem can be used to approximate credit losses. Second, the exposure on instrument i must be independent of the exposure on instrument j , $i \neq j$. This assumption can be slightly weakened; all that is required is that the exposure on instrument i be independent of the distribution of credit losses on the rest of the portfolio for all i . Third, the exposure on each instrument must be independent of the process generating counterparty default probabilities.

As argued above, this third assumption is likely to be false for typical derivative instruments. The second assumption is even more tenuous. Hull's version of this assumption will never be true as long as there is more than one instrument of a given type in the portfolio (say, at least two U.S. dollar interest-rate swaps). The weaker version of the assumption requires that the distribution of credit losses on the portfolio be independent of the distribution of any relevant financial variables. Given the third assumption, this weaker second assumption requires that the aggregate exposure of the portfolio must be independent of the underlying financial variables that determine the exposures on the individual instruments.

When these assumptions do not hold, the marginal effect of a derivative instrument on the upper bound of credit losses associated with a portfolio must be measured in the context of the portfolio. The marginal effect will be relatively high (low) when the exposure on the instrument is positively (negatively) correlated with the exposure on the rest of the portfolio; it will also be relatively high (low) when the exposure is positively (negatively) correlated with the number of defaulting counterparties.

The remainder of this discussion illustrates the difference between the true marginal impact of a given derivative instrument on a portfolio's credit losses and the estimated impact as measured in the typical manner. I calculate the 95th percentile upper bounds on credit losses for four derivative instruments (actually

both sides of two instruments). I first calculate these upper bounds in the standard way: first calculating the average 95th percentile upper bound on exposure, then multiplying by the cumulative probability of counterparty default. I then calculate these upper bounds in the context of a particular portfolio.

For this exercise, I arbitrarily choose a portfolio that consists of 100 derivative contracts with 100 counterparties. One-fourth of the contracts are paying-fixed on five-year U.S. dollar interest-rate swaps, one-fourth are receiving-fixed on the same swap, one-fourth are receiving dollars on five-year dollar/DM currency swaps, and the remainder are paying dollars on the same swap. All of these contracts have the same notional value and were initiated at the same time. I consider four different marginal contracts: each side of a five-year U.S. dollar interest-rate swap and each side of a five-year dollar/DM currency swap. The marginal contract is with yet another counterparty, so that the total portfolio (including the additional marginal contract) consists of 101 instruments with 101 counterparties.

This portfolio should be no means be interpreted as representative; I chose it because there is sufficient historical data on U.S. and German term structures, as well as the \$/DM exchange rate, so that the historical distribution of exposures can be used instead of a distribution from a Monte Carlo simulation. In practice, the 'portfolio' should be the entire portfolio of the firm, including loans and other sources of credit exposure, not just the portfolio of derivative instruments.

The unconditional distribution of credit losses was generated by Monte Carlo simulation using the joint behavior of financial variables and default probabilities used above in the calculations of expected exposures and credit losses. This simulation implicitly assumes that the counterparties are typical firms rated by Moody's. In other words, the counterparties are drawn from a well-diversified group. This implicit assumption will affect the distribution of credit losses that is calculated here, because this distribution depends, in part, on the correlations among different counterparties' defaults. The tails of the distribution of credit losses are fatter for counterparties that have highly correlated default probabilities. Because of industry-specific shocks, these correlations will be larger for firms within a given industry than they will be for firms in different industries. Actual dealer portfolios may be far less diversified than the universe of Moody's firms, and therefore may exhibit fatter tails than the distribution calculated here.

An 'observation' of credit losses on the portfolio is generated as follows. First, one of the 204 different five-year paths of financial variables and default probabilities was randomly chosen. For this time path there is a cumulative probability of default over five years. Next, 101 independent random draws were made, using this cumulative probability of default, to determine whether counterparty i , $i = 1, \dots, 101$, defaulted at some time along this path. For each counterparty that defaults, another random drawing is made to determine at what point along this path the default occurred. This random drawing uses the marginal probabilities of default in months one through sixty. For each 'observation', we record the total

Table 3

A comparison of different methods of calculating an upper bound on credit losses on certain derivative instruments

Contract	Average 95th percentile bound on exposure ^a	Implied credit loss ^b	Marginal effect on 95th percentile bound on portfolio's credit loss ^c
Five-year US\$ interest-rate swap			
Receive-fixed	9.848	0.055	0.149
Pay-fixed	8.937	0.050	0.000
Five-year US\$/DM currency swap			
Pay-\$	51.06	0.287	1.837
Receive-\$	27.05	0.152	0.149

Distributions of credit exposures are calculated with actual U.S. and German risk-free term structures and the spot \$/DM exchange rate, over 1971–1992. Default behavior over the same period is for firms originally rated 'A' by Moody's. Losses are calculated before any recovery in bankruptcy. All figures are in percent of notional principal.

^a The unweighted average of the 95th percentile bound at each month-end during the life of the instrument.

^b The average 95th percentile bound multiplied by an A-rated firm's five-year cumulative probability of default.

^c The actual marginal impact of each instrument on a given portfolio's 95th percentile bound on credit losses. This portfolio consists of 100 contracts, each with different counterparties. It is divided evenly into pay-fixed U.S. dollar five-year interest-rate swaps, receive-fixed U.S. dollar five-year interest-rate swaps, pay-dollar U.S./DM five-year currency swaps, and receive-dollar U.S./DM five-year currency swaps.

credit losses for both the original portfolio of 100 instruments and the new portfolio of 101 instruments. I make no attempt to discount the losses to reflect the times of the various defaults. This procedure was replicated 50 000 times to create the distribution of undiscounted credit losses prior to recovery in bankruptcy.

To calculate the impact of each of the four marginal derivative instruments, I first calculate the 95th percentile upper bound on aggregate credit losses for the original portfolio of 100 instruments. I then add the marginal derivative contract to the portfolio and recalculate the 95th percentile upper bound on aggregate credit losses. The change in this upper bound is the true marginal effect of this additional derivative contract.

Table 3 reports three numbers for each marginal instrument. The first column reports the average 95th percentile upper bound on exposure. The second column translates this number into a credit loss by multiplying by 0.00563 (the average cumulative default probability in this data). The third column reports the marginal increase in the portfolio's 95th percentile bound on credit losses given by the above Monte Carlo simulation.

The portfolio's total credit losses are largely associated with the counterparties

to the receive-dollar side of the currency swap (79% of all credit losses are the result of defaults by the counterparty on the other side of this swap). Therefore the actual marginal impact of another receive-dollar side is very large: 1.8% of the notional value of the swap. This impact is over six times larger than that implied by the 95th percentile bound on this instrument's exposure. By contrast, the actual marginal impact of the pay-dollar side of this currency swap (0.15% of notional) is slightly less than that implied by this instrument's bound on exposure.

Over the past twenty years, the value of the dollar and the level of U.S. interest rates have generally moved together. Therefore in this Monte Carlo simulation, credit losses on receive-fixed interest-rate swaps are positively correlated with losses on pay-dollar currency swaps, and hence with losses on the portfolio as a whole. The marginal effect of another receive-fixed interest-rate contract is therefore much larger (three times larger) than that implied by the contract's upper bound on exposure. By contrast, the marginal impact of another pay-fixed side of the interest-rate swap is zero. In other words, the credit losses on this contract were not realized at times when the total credit losses on the portfolio were extremely large.

4. Measuring credit losses over short horizons

Throughout the preceding discussion of credit risk, risk was measured over the remaining life of the contract or contracts. However, it is not clear that this horizon is the appropriate one. Derivative contracts need not be held until maturity; they can be terminated or assigned to other dealers. The industry is increasingly looking to these alternatives as mechanisms for managing credit risk.

Termination or assignment are tools available to regulators as well. For example, the Bank of New England had a derivatives book of \$30 billion in notional principal in January 1990. At that time the Bank was in considerable financial difficulty. Over the next year, the size of the book was reduced through terminations and assignments under the supervision of banking regulators. The book had a remaining notional principal of \$6.7 billion by January 1991, at which time the bank was declared insolvent. Drexel's derivatives book was also substantially unwound prior to the bankruptcy of its derivatives subsidiary. A good summary of these events is Cunningham (undated).

In this section we consider how to measure credit losses over a time horizon shorter than the instrument's life. Before we get into the details of this procedure, however, we should think about the proper time period over which credit losses should be measured. We do so from the perspective of a bank regulator. The Risk-Based Capital Standards require banks to hold capital against current exposures on derivative instruments and additional capital to cover potential future increases in exposures. What risks should these 'add-ons' be designed to cover?

One line of reasoning is that add-ons provide a buffer stock of capital to cover

potential increases in credit exposures between the times that banks measure their current exposures. This reasoning is implicit in much of the banking industry's discussion of add-ons (see, e.g., Bankers Trust, 1987). This approach implies that if banks recalculated their current exposures very frequently (say, every day), they would need very small add-ons, or perhaps none at all. A problem with this approach is that it presumes that when banks discover that their current exposures have increased, they can simply raise additional capital to cover their new levels of current exposures. Of course, this presumption also implies that bank failures would never happen as long as bankers were constantly aware of the market value of their asset portfolio. No one would disagree with the idea that bankers should be aware of their portfolio's market value, but the notion that such knowledge is sufficient to ward off bank failures is, at best, problematic.

An alternative view of add-ons is that they are designed to cover potential increases in credit exposures during the time it takes for regulators to both detect a situation in which a bank will have insufficient capital to cover future increases in credit exposures and to respond to this situation by unwinding a bank's derivatives portfolio. The amount of time required will vary from bank to bank, depending on the size and complexity of the bank's portfolio as well as market conditions at the time the bank gets into difficulty.

This paper will not pretend to determine the length of time needed to detect a problem and unwind a bank's portfolio. We leave that for further consideration. In the remainder of this section we consider how credit risks should be measured for a derivative instrument or instruments over a given period of time shorter than the maturity of the contract(s).

When we focus on horizons shorter than the life of the instrument, the second definition of credit loss mentioned in Section 2.2 is the appropriate definition: A credit loss is the loss incurred when the value of a derivative instrument falls owing to changes in counterparty credit quality.⁴ (Conversely, there are also credit gains: increases in value owing to an increase in credit quality.) Dealers in derivative markets report that when contracts are unwound, their 'resale' (i.e. assignment) value depends on the health of the original counterparty at the time of assignment. This is likely to be true even if contracts are terminated instead of assigned. Sophisticated counterparties recognize the effect that their own credit-worthiness has on the mark-to-market value of a derivative contract. For example, when a securities firm experienced difficulties several years ago, some of its counterparties raised the possibility of terminating some of their derivatives contracts. The firm reportedly indicated that it would consider terminations as long

⁴ An alternative definition is that a credit loss is incurred whenever the difference between the value of the default-risky instrument and the value of an identical default-free instrument increases. With this definition, credit losses can be incurred without any change in counterparty credit quality because an increase in the value of the obligated cash flows will be discounted by the probability of firm survival.

as its derivative obligations were discounted at the same rate as the market's discount on its senior debt.

There are very few nonproprietary empirical estimates of distributions of credit losses on derivative instruments over fixed time horizons, and they all (explicitly or implicitly) assume away the complications introduced by time-varying credit quality. Bankers Trust (1987) and Hendricks and Barker (1992) estimate distributions of credit exposures over some arbitrary time horizon. The expected credit loss over that horizon is then the expected credit exposure over that horizon times the probability of counterparty default over the same horizon. Hull (1989a) justifies this approach by assuming that counterparties have a constant probability of default over every time interval Δt . Therefore there are no perceived variations in credit quality to affect the value of the derivative instruments. This assumption is at odds with most of the valuation models discussed in Section 2.1 as well as reality. Firms' perceived credit quality varies over time, and firms with publicly traded debt rarely default without prior downgrades (Altman and Kao, 1991; Carty and Fons, 1993; Fons et al., 1994; Brand et al., 1994; and Jarrow et al., 1994).

Without an assumption like Hull's, estimating distributions of credit losses is very difficult. The starting point for such estimation is a valuation model such as those described in Section 2.1. These models determine mark-to-market values of derivative instruments as functions of the joint stochastic behavior, under the equivalent martingale measure, of counterparty credit quality and the variables that determine obligated cash flows. To estimate distributions of credit losses, the variability over time of these mark-to-market values owing to variations in counterparty credit quality must be estimated.

5. Conclusion

This paper reviews the current techniques used to measure the credit risks associated with derivative instruments over the lives of these instruments. It identifies major problems with these techniques. All of these problems have been previously noted, but then largely ignored. The contribution of this paper is to document that these problems lead to substantial mismeasurement of credit risks for typical derivative instruments and counterparties.

The first problem is that the stochastic models used in Monte Carlo simulations are too often accepted at face value. For example, the parameters in these models are estimated, but the simulations ignore the fact that they are estimated, and instead treat the parameters as known. In addition, these models are inherently simplistic, tending to ignore important features of the behavior of the underlying financial variables. Addressing this problem requires a greater focus on econometric estimation and less on simple calibration.

The second problem is that the correlation between exposure on a given derivative instrument and the counterparty's default probability is largely ignored.

Defaults are primarily driven by the business cycle, which also drives variations in the financial variables on which most derivative instruments are priced. Research efforts must focus on the joint determinants of variations in firms' credit quality and variations in financial variables, such as interest rates and equity prices.

Finally, upper bounds on credit losses associated with given derivative instruments are calculated without regard to the total portfolio of the dealer or end-user. However, upper bounds on credit losses cannot be computed meaningfully in isolation. An upper bound on credit loss for an individual instrument should be the marginal impact of the instrument on the upper bound of credit losses for the entire portfolio.

Most of this paper is concerned with measuring credit risk over the life of a derivative instrument. However, since contracts can be terminated prior to their maturity, a shorter horizon may be more relevant. If so, we need to change our definition of credit loss. Losses are accrued whenever a counterparty's credit quality is perceived to fall: A decline in credit quality results in a worse assignment or termination value for the instrument. Techniques for measuring credit risk must change accordingly.

A tremendous amount of work is currently being done by researchers, both inside and outside the derivatives industry, to increase the speed of credit risk measurement. Such increases in speed are necessary in order for large derivative dealers to assess the credit risks they face in a timely manner. The results in this paper suggest that similar efforts are required to make these timely measurements accurate as well.

Acknowledgements

I thank Peter Abken, Robert Butler, John Kambhu, Patrick Parkinson, Patricia White, and two anonymous referees for helpful comments. Remaining errors are my own. The analysis and conclusions of this paper are those of the author and do not indicate concurrence by other members of the research staff, by the Board of Governors, or by the Federal Reserve Banks.

References

- Abken, P., 1993, Valuation of default-risky interest-rate swaps, *Advances in Futures and Options Research* 6, 93–116.
- Altman, E.I., 1989, Measuring corporate bond mortality and performance, *Journal of Finance* 44, 909–922.
- Altman, E.I., R. Haldeman and P. Narayanan, 1977, ZETA analysis: A new model to identify bankruptcy risk of corporations, *Journal of Banking and Finance* 1, 29–54.
- Altman, E.I. and D.L. Kao, 1991, Examining and modeling corporate bond rating drift, Working paper (New York University Salomon Center, New York, NY).

- Arak, M., L. Goodman and A. Rones, 1986, Credit lines for new instruments: Swaps, over-the-counter options, forwards and floor–ceiling agreements, in: Federal Reserve Bank of Chicago, Proceedings of a conference on bank structure and competition, Vol. 23 (Federal Reserve Bank of Chicago, Chicago) 437–456.
- Bank of England and Federal Reserve Board Staff, 1987, Potential credit exposure on interest rate and foreign exchange rate related instruments, Working paper (Federal Reserve Board, Washington, DC).
- Bankers Trust, 1987, An appreciation of the analysis of credit exposure on single-currency (U.S. dollar) fixed–floating interest rate swaps in 'Treatment of interest rate and exchange rate contracts in the risk asset ratio' and 'Potential credit exposure on interest rate and exchange rate related instruments', Working paper (Bankers Trust Money Market Center, New York, NY).
- Belton, T., 1987, Credit risk in interest rate swaps, Working paper (Federal Reserve Board, Washington, DC).
- Black, F. and J.C. Cox, 1976, Valuing corporate securities: Some effects of bond indenture provisions, *Journal of Finance* 31, 351–367.
- Black, F. and M. Scholes, 1973, The pricing of options and corporate liabilities, *Journal of Political Economy* 81, 637–654.
- Bond, I., G. Murphy and G. Robinson, 1994, Potential credit exposure on interest rate swaps, Working paper (Bank of England, London, UK).
- Brand, L., T. Kotto and R. Bahar, 1994, 1993 corporate default, rating transition study results, Standard & Poor's CreditReview (May), 1–12.
- Carty, L. and J. Fons, 1993, Measuring changes in corporate credit quality, Working paper (Moody's, New York, NY).
- Chen, R. and L. Scott, 1993, Maximum likelihood estimation for a multifactor equilibrium model of the term structure of interest rates, *Journal of Fixed Income* 3, 14–32.
- Cooper, I. and A. Mello, 1991, The default risk of swaps, *Journal of Finance* 46, 597–620.
- Cox, J.C., J. Ingersoll and S. Ross, 1985, A theory of the term structure of interest rates, *Econometrica* 53, 385–407.
- Cunningham, D.P., undated, Swaps and termination events: Legal and business considerations, Working paper (Cravath, Swaine and Moore, New York, NY).
- Duffee, G., 1992, The replacement cost of equity index swaps, Working paper (Federal Reserve Board, Washington, DC).
- Duffie, D. and M. Huang, 1994, Swap rates and credit quality, Working paper (Stanford Graduate School of Business, Stanford, CA).
- Duffie, D. and K.J. Singleton, 1994, Econometric modeling of term structures of defaultable bonds, Working paper (Stanford Graduate School of Business, Stanford, CA).
- Estrella, A. and G. Hardouvelis, 1991, The term structure as a predictor of real economic activity, *Journal of Finance* 46, 555–576.
- Ferron, M. and G. Handjinicolaou, 1987, Understanding swap credit risk: The simulation approach, *Journal of International Security Markets* 1, 135–148.
- Fons, J., L. Carty and J. Kaufman, 1994, Corporate bond defaults and default rates, Moody's Special Report, January.
- Franks, J.R. and W.N. Torous, 1994, A comparison of financial recontracting in distressed exchanges and Chapter 11 reorganizations, *Journal of Financial Economics* 35, 349–370.
- Geske, R., 1977, The valuation of corporate liabilities as compound options, *Journal of Financial and Quantitative Analysis* 12, 541–552.
- Gilberti, D., M. Mentini and P. Scabellone, 1993, The valuation of credit risk in swaps: Methodological issues and empirical results, *Journal of Fixed Income* 2, 24–36.
- Gluck, J. and B. Clarkson, 1993, Moody's approach to evaluating derivative products subsidiaries, Moody's Special Comment (October).
- Group of Thirty, 1993, Derivatives: Practices and principles (Group of Thirty, Washington).

- Hendricks, D. and D. Barker, 1992, Monte Carlo simulations for interest rate swaps, Working paper, Federal Reserve Bank of New York, New York, NY.
- Ho, T. and S. Lee, 1986, Term structure movements and pricing interest rate contingent claims, *Journal of Finance* 41, 1011–1029.
- Hull, J., 1989a, Assessing credit risk in a financial institution's off-balance sheet commitments, *Journal of Financial and Quantitative Analysis* 24, 489–501.
- Hull, J., 1989b, An analysis of the credit risks in interest rate swaps and currency swaps, *Recent Developments in International Banking and Finance* 3, 109–130.
- Hull, J. and A. White, 1992, The impact of default risk on the prices of options and other derivative securities, Working paper (University of Toronto, Toronto, Ont.).
- Jarrow, R., D. Lando and S. Turnbull, 1994, A Markov model for the term structure of credit risk spreads, Working paper (Cornell University, Ithaca, NY).
- Lando, D., 1994, Three essays on contingent claims pricing, Ph.D. dissertation (Cornell University, Ithaca, NY).
- Longstaff, F. and E. Schwartz, 1994, A simple approach to valuing risky fixed and floating rate debt and determining swap spreads, Working paper (Graduate School of Management, UCLA, Los Angeles, CA).
- Lucas, D., 1994, Quantitative measures of credit risk, Working paper (TMG Financial Products).
- Madan, D.B. and H. Unal, 1994, Pricing the risks of default, Working paper (Wharton, Philadelphia, PA).
- Merton, R.C., 1974, On the pricing of corporate debt: The risk structure of interest rates, *Journal of Finance* 29, 449–470.
- Muffet, M., 1987, Modeling credit exposures on swaps, in: Federal Reserve Bank of Chicago, *Proceedings of a Conference on Bank Structure and Competition*, Vol. 23 (Federal Reserve Bank of Chicago, Chicago) 473–496.
- Neal, K. and K. Simmons, 1988, Interest rate swaps, currency swaps and credit risk, *Issues in Bank Regulation* 11, 26–29.
- Standard & Poor's, 1992, Banks face manageable risks in derivative businesses, *Standard & Poor's CreditReview* (November), 46–51.
- Shimko, D.C., N. Tejima and D.R. van Deventer, 1993, The pricing of risky debt when interest rates are stochastic, *Journal of Fixed Income* 3, 58–65.
- Sorensen, E.H. and T.F. Bollier, 1994, Pricing swap default risk, *Financial Analysts Journal* 50, 23–33.
- Vasicek, O.A., 1977, An equilibrium characterization of the term structure, *Journal of Financial Economics* 5, 177–188.
- Whittaker, J.G., 1987, Interest rate swaps: Risk and regulation, Federal Reserve Bank of Kansas City *Economic Review* 72, 3–13.