NEW YORK UNIVERSITY
NYU STERN
LEONARD N. STERN SCHOOL OF BUSINESS

Department of Finance
Working Paper Series 1998

FIN-98-033

Derivatives Risks, Old and New

Stephen Figlewski

Dec. 7, 1997

This working paper series has been generously supported by a grant from

CDC Investment Management Corporation

Version of Dec. 7, 1997

Comments Welcome

Derivatives Risks, Old and New

by

Stephen Figlewski

Professor of Finance

New York University Stern School of Business
44 West 4th St., Suite 9-160
New York, NY 10012-1126

Phone: 212-998-0712
FAX: 212-995-4220

E-mail: sfiglews@stern.nyu.edu

I would like to thank T. Clifton Green and Sari Carp for invaluable assistance in preparing this paper. I am grateful for valuable discussions and comments on earlier drafts provided by Kenneth Froewiss, Robert Litan, Leslie Rahl, Alexander Reisz, Anthony Santomero, Charles Smithson, René Stulz and seminar participants at NYU and at the Brookings-Wharton Conference.

Abstract

There has been much discussion of risks tied to trading in derivatives, with some well-informed objective observers arguing that derivatives risks are not significantly greater or different from those associated with traditional financial instruments. Financial risks are often broken down into market risk, credit risk, operational risk and legal risk. We review the standard classification and observe that while derivatives are exposed to these types of risk, they are manifested quite differently in derivatives than in traditional securities. We then consider a "new" type of risk that is particularly important for derivatives: model risk. Derivatives trading depends heavily on the use of theoretical valuation models, but these are susceptible to error from incorrect assumptions about the underlying asset price process, estimation error on volatility and other inputs that must be forecasted, errors in implementing the theoretical models, and differences between market prices and theoretical values. Empirical evidence drawn from several important asset markets shows that model error can be quite large and can be expected to lead to significant risk in derivatives pricing and risk management.

Introduction

Although derivative instruments have been traded for a long time, the enormous growth in the volume and variety of futures, options, swaps, and more exotic types of contracts in recent years has been without precedent. Concern about the risks of trading in these instruments is also not new, but it too has grown along with the markets. In the last couple of years, a series of widely publicized losses related to derivatives activities has focused public attention (once again) on derivatives risks.

The tone of the discussion has evolved, however, from calls to suppress trading in markets that are asserted to be too speculative, like the onion futures market that was closed down by act of Congress in 1958, to a more constructive recognition that these instruments are now a permanent feature of our financial markets and that it is necessary to find ways to assess and manage the risks they entail.

Any objective assessment of financial derivatives has to conclude that these markets have contributed greatly to our ability to manage economic and financial risk. Derivatives are invaluable in separating the bearing of risk from the natural exposure to that risk that results from one's ownership of risky assets or from one's economic position generally. For example, derivatives markets allow inventories of commodities and securities to be carried without the necessity of also bearing the risk of price fluctuation. The traditional role of futures markets as vehicles for hedging commodity risk has been extended to many kinds of financial instruments that entail much greater aggregate exposure to price risk than do traditional commodities. Derivatives make it possible for firms to obtain financing wherever and however it is cheapest and to transform the resulting debt into the form that is desired. Today, a firm wishing to borrow dollars at a fixed interest rate for 5 years may find it cheaper to borrow Japanese yen at a floating rate and to use currency forwards and an interest rate swap to offset the exchange rate risk and turn the floating rate liability into one with a fixed interest rate.

Through derivatives, major classes of risk that in the past were mostly borne by specialized financial institutions, with limited risk bearing capacity, can now be shared more broadly. For example, derivatives based on mortgages have made it possible for home buyers to acquire funds from the bond market rather than having to rely on the ability of savings and loans and similar financial institutions to attract deposits. Recent innovations in derivatives based on catastrophic risks like hurricanes and earthquakes are beginning to make it possible for insurance companies to share risk exposure more broadly with outside investors. Derivatives with option features allow investors to restructure risk exposures to provide preferred patterns. For the public, this often means allowing an investor to limit the risk of a loss from an adverse price change without eliminating profits from a favorable market move.

But along with the benefits of powerful new tools for managing risks and for creating preferred returns patterns that derivatives provide have also come what often appear to be substantial new risks tied to the derivative instruments themselves. Derivatives are distinctly more complicated than stocks, bonds, loans, bank accounts, and other traditional financial instruments. In this paper we will refer to these traditional classes of securities as "fundamental" securities. Although they are comfortable with fundamental securities, the public at large have relatively little understanding of derivatives, and the large derivatives losses experienced by major corporations and financial institutions in recent years suggest that even sophisticated investors are capable of making big mistakes about derivatives.

The large and growing importance of derivatives to our financial system, coupled with the perception that they entail significant risk that not all investors are fully prepared to deal with, has prompted several high-level groups to study the issues of risk in derivatives, with an eye to promoting general principles and effective practices for managing it, and also in some cases, with the objective of instituting more formal regulatory policies for reporting and limiting risk exposures for banks and other regulated financial institutions. These include a major study by the Group of Thirty, with follow-up surveys of industry practice; a study conducted by the U.S. General Accounting Office (GAO) published in 1994; studies done by the Bank for International Settlements (BIS) which developed principles that have been embodied in the bank capital standards recently established for all banks in the European Community (EC); a set of proposed risk standards for institutional investors developed by the private Risk Standards Working Group; and continuing attention by the Financial Accounting Standards Board (FASB) to the difficult accounting issues raised by derivatives.

With regard to derivatives risks, a common theme in these studies, as expressed by the Chairman of the Group of Thirty, Paul Volcker, in the Foreword to their report on <u>Derivatives: Practices and Principles</u> (p. *i*) [1993] is,

> "The general attitude of the Study towards regulation is plain: derivatives by their nature do not introduce risks of a fundamentally different kind or of a greater scale than those already present in the financial markets. Hence, systemic risks are not appreciably aggravated, and supervisory concerns can be addressed within present regulatory structures and approaches."

The Group of Thirty study offered 24 recommendations to market participants and to regulators regarding the management of derivatives-related risks. Several of the other studies also made explicit recommendations of a similar nature.

The principles of derivatives risk management contained in these recommendations are clear, comprehensive, and likely to be very effective if put into general practice. Most of the avoidable derivatives losses that have created headlines in the news have not occurred because the recommendations were inadequate but because they were not followed.

However, while Volcker's statement that derivatives do not introduce risks of a fundamentally different kind is basically correct, it does not capture an important aspect of derivatives, which is that the kinds of risks present in derivatives are the same as in fundamental securities, but the ways these risks are manifested are often significantly different, and new ways of understanding traditional sources of risk as they apply to derivatives are required.

Moreover, there is one important type of risk that is essentially new with derivatives: model risk. Derivatives are complex financial instruments but the theory of how they should be priced and how they can be expected to respond to changes in market conditions is well developed. The theoretical principles are incorporated into mathematical models, and virtually all serious derivatives traders have access to computer implementations of these models and depend on them in trading derivatives and in assessing and managing risk exposures.

But although derivatives models may be rigorously derived from accepted theoretical principles, and involve complex equations and daunting mathematics understood only by "rocket scientists," they remain only models of reality. To the extent that the real world differs from the models, reliance on them will lead to risk exposure due to model inaccuracy. The problem is compounded by the fact that derivatives models require the user to input a number of parameters, including some that are not directly observable, like the volatility of the underlying asset. Volatility can be forecasted using standard and not-so-standard techniques, but there will necessarily be forecast errors that add to model risk.

Some problems with the models are generally known, but no good solution is available. An example is the widespread use of the lognormal probability distribution with constant volatility for security returns, as in the Black-Scholes (BS) option pricing model, even though it is well-known that volatility varies over time, and actual returns in virtually every market that has been examined have "fat tails" (that is, the actual probability of a large price change is greater than the model allows for). Often the existence of these problems is known, but their magnitude, in terms of their impact on the risk exposure of a given investment position or trading strategy, is not.

It is the unique character of derivatives risks, and particularly model risk, that will be the main focus of this paper. We begin, in Section II, with an overview of the particular ways in which the traditional sources of financial risk impact derivatives. Section III then discusses model risk in pricing and hedging derivatives and presents some estimates of its magnitude. The final section presents concluding comments.

II. Traditional Sources of Risk in Derivatives

As with other financial instruments, derivatives risks can be divided into credit risk, market risk, legal risk and operational risk:

Market risk: the risk that movements in financial market prices will impair a firm's financial condition due to its positions in derivatives.

Credit risk: the risk (broadly defined) that the counterparty to a derivatives contract will fail to fulfill its contracted obligations.

Operational risk: the risk of derivatives-related losses from deficient internal controls or information systems.

Legal risk: the risk that derivatives contracts will not be legally enforceable.

The studies cited above show plainly that the unique characteristics of derivative instruments with regard to these four kinds of risks are actually well-understood in both the private and the public sectors.

The guidelines and procedures for managing derivatives risks that they propose are sensible, reasonable, and likely to be effective if implemented generally. It is safe to say that widespread adoption of the precepts they propose would greatly reduce the incidence of avoidable derivatives losses, including the most widely publicized events of the last several years, such as those by Barings, Metallgesellschaft, and Procter & Gamble.

One other class of risk that is widely discussed with respect to derivatives is "systemic" risk, that is, the risk that an event originating in a derivatives market could spread to other markets and precipitate a general financial crisis. Like the perennial question of whether speculation stabilizes or destabilizes a financial market, about which much has been said without producing general agreement, systemic risk is an important issue that will not be resolved by a single argument. In the last subsection, I offer a brief argument that systemic risk from derivatives activities should be much smaller than that associated with fundamental securities, because derivatives are a zero sum game, and can not create or destroy aggregate wealth.

II.1 Market Risk in Derivatives

Market risk is the risk that price changes in the financial markets will cause a loss on a securities position. Evaluating exposure to market risk is usually straightforward for fundamental assets. For instance, if the stock market falls, a financial institution holding a portfolio of equities can expect to experience a loss on that portfolio that will be directly related to the size of the market move. The predictions of modern portfolio theory on this issue are well-established and empirically supported: a random fluctuation in the broad stock market can be expect to affect a particular stock portfolio in proportion to the portfolio's beta coefficient. If the market moves x%, a portfolio with a beta of 1.0 is expected also to move x%, while one with a beta of 1.5 should move by 1.5x %. Firms forecast stock betas by statistical estimation using past returns for the stock and a market index, and these betas aggregate into beta values for portfolios. The proportional relationship makes estimating market risk exposure for a portfolio relatively easy. Residual risk, i.e., changes in portfolio value that are independent of the broad stock market, is also easily estimated from the same kind of statistical analysis.

Similar approaches to market risk assessment apply to other classes of fundamental assets. Bond prices go down when interest rates rise, according to a fixed, though nonlinear relationship. Given a measure of variability in market yields, the market risk on a bond portfolio is easily computed. The value of a bank's position in a foreign currency will vary directly with the exchange rate, and so on.

Market risk exposure for derivatives positions is basically similar, and yet the results are often quite different in practice. For example, a financial institution that has written a call option based on the stock market index also experiences gains and losses in the value of its position as the market fluctuates, but the relationship between the market move and the change in a derivative's value is typically significantly more complex than the risk exposure of a portfolio of stocks. First, the direction of the relationship is as likely to be negative as positive; e.g., for a short call position, it is a rising stock market that will cause a loss. Second, the relationship is nonlinear, and in some cases highly nonlinear. The loss caused by a large market move can be proportionally much larger than one for a small move: For example, a 1 percent move in the market might cause a 5 percent loss on an option position, while a 2 percent market move would produce a 20 percent loss.

Third, due to higher leverage, risk exposure relative to the dollar value of a position is normally much greater for derivative instruments than for fundamental assets. A purchased option that ends up out of the money experiences a loss of 100% of the purchase price, even though the underlying asset may have moved only a little, or not at all. Moreover, every option contract has both a long side and a short side, and for an option that ends up in the money, the option writer's (i.e., the short's) potential loss is essentially unlimited; it can easily far exceed 100% of the initial price of the option. Finally, fundamental assets have fairly straightforward relationships with market risk factors, like that connecting the value of a stock portfolio to the return on the market index. But derivatives market risk exposure is more complex, and evaluating and managing it typically requires the use of mathematical valuation models. These models incorporate factors like volatility that reflect new, specifically derivatives-related, forms of market risk.

Because of the particular properties of market exposure for derivatives, dealer firms normally hedge their positions. Quantitative valuation models are used extensively for pricing derivatives, but their most important use is actually for managing derivatives risks.[1] The change in a

[1] Traders employ valuation models actively to price derivatives, of course, but they are typically first tuned to current market prices by the use of implied volatility as the model's volatility parameter. Implied volatility is the value that sets model prices equal to observed

derivative's value that is caused by a given (small) change in the market value for the underlying asset is the "delta." To insulate a derivatives position against market risk, a hedger takes a position in some combination of the underlying asset and other derivatives based on it that has a delta equal in size and opposite in sign from the position to be hedged. The resulting hedged position is said to be "delta neutral," meaning that its value will not be affected either up or down by a small change in the underlying asset. But in practice, delta hedging is just the beginning of risk management for derivatives.

To illustrate how market risk actually impacts a typical derivatives dealer firm, let us take the example of a bank that writes 3 month European call options on the Japanese yen. We assume the spot exchange rate (S) is 90.00 U.S. cents per 100 yen and the call is struck at the money, i.e., the strike price X=90. For convenience in expressing the dollar values involved, we assume one call option is for 100 yen. We will analyze the risk exposures related to this position using the standard Garman-Kohlhagen [1983] currency option pricing model. This is a variant of the Black-Scholes [1973] model, modified slightly to apply to exchange rates. The equation is shown in the Appendix. To use the model, we must also specify both U.S. and foreign interest rates (r and r_{YEN}, respectively), and the volatility of the exchange rate (σ). We assume these values are: r = 6%, r_{YEN}=2%, and σ=12%, which are representative values for the recent past.

Putting these parameter values into the model produces the estimates shown in the first part of Table 1. In addition to the model value for the call price of 2.57, a variety of other "Greek letters" are computed. These are computed from partial derivatives of the option value function and indicate the sensitivity of the model value to changes in the other input parameters. Delta is 0.57, meaning that for a small change in the exchange rate, the value of the call on 100 yen will change by about the same amount as a long spot position in 57 yen. However, the option value is a nonlinear function of the exchange rate, so the delta will change as the rate moves. The change in delta per unit change in the spot rate is given by the gamma. A large gamma means that the sensitivity of the option price to the underlying changes sharply as the market moves away from its current level. In this case, a one point rise in the exchange rate will increase the delta by about 0.73.

The change in option value is smaller than the change in the underlying but it is a larger percentage. The leverage ratio or elasticity of the option price is sometimes referred to as lambda. The value of 20.1 indicates that a 1 percent change in the exchange rate would produce about a 20 percent change in the option value.

---

market prices, so in effect "model values" come largely from the market. Theoretical pricing models are needed most for predicting how a derivative's market price will behave as the underlying asset price and other parameters change, so that the resulting risk exposures can be managed.

Other factors that influence option value and need to be taken into account in managing derivatives risk include time decay, measured by theta, sensitivity to changes in volatility, measured by vega[2], and sensitivity to interest rates, measured by rho. An option's value is partly a function of its time to expiration. Theta is the change in value if the underlying price is constant as one day elapses. For this option, if the exchange rate does not move for a day, it should lose value by about 0.017 cents per 100 yen.

Volatility is the one input parameter to a theoretical valuation model that is hardest to judge accurately. Unlike the option's strike and maturity which are fixed, or the asset price and interest rate that can be observed directly in the market, volatility must be forecasted. There are several basic approaches for obtaining a value for the volatility input, but all are subject to error. Moreover, the implied volatility (which is often thought of as the "market's" volatility forecast, as embodied in current market option prices) can vary substantially from one day to the next. Thus vega is an important risk measure in assessing the exposure of a derivatives position to factors that affect its market value. For this option, a rise in σ from 0.12 to 0.13 should raise its value about 0.174.

With rho of 0.001, sensitivity to interest rates (holding the underlying constant) for this option is very low. Longer term instruments have larger rho values.

To illustrate how a derivatives dealer might evaluate and manage the market risk involved in writing options, let us consider a dealer who has sold $100 worth of these 3 month yen calls. (At the model call price, this amounts to writing $100/ $0.0257 = 3894 calls, each based on 100 yen.) We will use the pricing model to simulate the profit or loss on this position over one day for a range of possible exchange rates, and then to examine how the risk exposure might be hedged. The results of the analysis for the whole position are shown in the lower portion of Table 1 and are plotted in Figures 1 and 2. With an initial position value of 100, the profit and loss figures can be interpreted as percentage returns.

First consider the behavior of the short call position with no hedging. This is plotted as the solid line in Figure 1. If the spot exchange rate does not change, the position value increases by less than 1 percent, due to the effect of time decay on the value of the calls that were sold. A drop in the spot rate produces a gain for the position, which is capped at 100 since the option value can not fall below zero. If the exchange rate rises, however, losses accrue at an accelerating rate, such that a 5 point rise produces a loss almost twice as large as the gain that would result from a 5 point drop. There is no cap on the possible loss.

---

[2] Vega, of course, is not a Greek letter, though it sounds like one and it starts with v, suggesting volatility. Purists may refer to an option's partial derivative with respect to volatility as kappa.

The position's directional market risk is due to the fact that the option delta is positive, meaning writing the option produces a negative delta. For a small change in the exchange rate, writing one contract is like taking a short position in 57 yen. Since the public much prefers buying options to selling them, derivatives market makers must be option writers on average. But they do not want to assume a large directional market risk to do so. The first step in managing derivatives risk exposure is to set up a delta neutral hedge. The textbook procedure is to take an opposite position in delta units of the underlying; in this case, the dealer would buy 57 yen per written call contract.

The market exposure of this delta neutral position is indicated by the curve with hollow boxes in Figure 1 and in Table 1, by the Delta Hedge line with volatility 0.12. Visibly, the position value is now much less affected by a small change in the spot rate and the sign of the result no longer depends on market direction. Delta hedging greatly reduces market risk exposure for the dealer, but there is still a problem. Due to the curvature of the option value function, a delta neutral hedge does not remain delta neutral as the underlying moves. A hedged position in which options are held short will lose money if a large price move occurs in either direction (and the position gains value from time decay if the underlying does not move). This is market risk of a new kind, that affects derivatives but not their underlying assets. Since market makers will typically hedge their positions, they do bear significant market risk, but it generally has quite a different character from that borne by investors in traditional securities.

The extent of the curvature in the hedged position is measured by its gamma, and gamma risk (also known as convexity risk) is a significant concern for a derivatives dealer, and for anyone who writes options. Dealers quickly understood this, and found ways to manage gamma risk along with delta. To offset the effect of curvature in the options that are sold, one can purchase options whose aggregate (positive) gamma is the same size as the position's negative gamma. Any options based on the same underlying can be used; here we consider buying 3-month 95 strike calls. Given the gamma of those options, it is necessary to buy 1.27 95 strike calls for each 90 strike call that is written. The resulting delta of the combined options position is -0.27, which is then hedged by buying 27 yen per 90 strike call contract being hedged.

The risk exposure of this delta-gamma hedge is illustrated by the curve marked with triangles in Figure 1. The effect of hedging both "Greek letter risks" is striking. The range over which this position has little market risk exposure has widened considerably: position value changes less than 10 percent for a 1 day exchange rate movement of 5 points either up or down (assuming nothing else, like volatility, changes). Again, we see that standard risk management procedures significantly transform the market risk exposure of derivatives dealers.

However, delta-gamma hedging does not eliminate all sources of derivatives risk. Both option value and the Greek letter risk exposures are functions of all of the input parameters to the valuation equation. Changes in volatility (implied volatility, to be precise, since it is variation in market option prices that is being hedged) will introduce risk into the hedge. To show the impact of volatility changes, we consider position values for three different volatility levels, the original value $\sigma = 0.12$, plus lower and higher values of $\sigma = 0.08$ and $\sigma = 0.16$. Figure 2 plots the effect on the delta-gamma hedge, and Table 1 shows results for both hedged positions.

The response of these positions to volatility risk is sizeable, and also complex. A change in volatility causes the market values of the options in the position to deviate from the values that were projected when the hedge was set up. Moreover, the pattern of market risk exposure is altered. For example, at $\sigma = 0.12$ the delta-gamma hedge gains value as the exchange rate rises, but with $\sigma = 0.08$, it first loses value for a moderate rise but then begins to gain when the yen gets above about 97. Notice that the basic delta hedge is more affected by volatility changes than the delta-gamma hedge when the spot rate is constant, because the latter contains options with largely offsetting vegas.

It should be no surprise that sophisticated derivatives dealer firms also try to hedge vega risk along with the other kinds. Moreover, sensitivities beyond these basic ones are scrutinized as well, like the change in delta when time elapses ("charm") and numerous others.

Two general principles should be clear from this discussion. First, while market risk is very important to derivatives dealers, the variety of risk management strategies that they routinely employ completely transform its impact on them. Derivatives market risk is significantly different from what investors in fundamental assets experience as market risk. Second, all of the derivatives risk assessment and risk management strategies we have discussed depend heavily on quantitative theoretical models. Plainly, if these models are wrong, derivatives dealers will be exposed to market risk in different amounts and in different ways from what they are expecting.

It is fair to say that for non-dealer investors and users of derivatives, the most important new form of market risk they bear related to derivatives is convexity, or gamma risk. The profile of returns for the written call options in Table 1 clearly illustrates the impact of convexity on an option position. Market risk exposure when options are written may be relatively small for small market moves, but it can accelerate rapidly for large moves. Serious derivatives-related losses occur when a trader is short options and the market makes a substantially larger move than he believes is possible.

Convexity risk can be insidious: a trader may hit upon a trading strategy that in one way or another amounts to writing deep out of the money options. A small option premium is earned each time and almost always the written options end up expiring out of the money. An incautious trader can come to feel that the strategy is a safe way to make money. But when a very large market move does cause the options to go into the money, the loss to the writer will typically be many times the premium that was received. Moreover, as we will discuss below, the probability of a very large market move is considerably greater in practice than would be expected under the lognormal distribution that is generally assumed for security prices. A recent example of this occurred as a result of the major stock market drop on October 27, 1997, when a well-established hedge fund run by Victor Niederhoffer was bankrupted. They had been writing deep out of the money put options on the Standard and Poors 500 stock index and collecting small amounts of option premium. When the stock market fell by more than 7.0 percent in one day--a virtual impossibility under the lognormal probability distribution assumed by the most prevalent option pricing models--the entire equity of the fund was lost.[3]

Some derivative instruments can have very large gammas under certain circumstances, especially those like barrier options that specify a price boundary at which the derivative's value undergoes a sharp change over a small price interval for the underlying. These derivatives are very hard to hedge effectively, which means that no good alternative may exist to simply bearing the convexity risk.

The fact that derivatives are a "zero sum game" has important implications for risk management. Since every contract must have both a long and a short side, if one trader makes a dollar profit, the counterparty on the other side loses a dollar. This principle extends to risk management: Derivatives shift risk, they do not eliminate it. If one dealer hedges her gamma risk by buying options, as in the above example, then someone else has to take on that risk. If the public prefers to buy options, dealers must on balance be short options. This means that while any one of them may hedge her gamma risk, as a group they must bear all of the risk that comes from taking the other side of the public's trades.

Since the effect of convexity is that delta changes as the underlying moves, the standard way to manage gamma risk is with a dynamic hedge, that alters the amount held of the hedging instruments as the delta changes. This is the heart of the arbitrage strategy Black and Scholes assumed in deriving the option pricing model. However, it has two significant shortcomings in practice that make it less than a perfect system for managing convexity risk. First, it is not actually possible to rebalance the hedge continuously, so the strategy can not eliminate all of the risk. A sharp change in the underlying while the market is closed overnight, for example, will still lead to a loss on the delta neutral hedge of a written option, as Figure 1 shows. Second, the way the hedge must work is that the underlying is sold as the market goes down and purchased as the market goes up. Thus hedge rebalancing generates positive feedback trading in the underlying. It becomes expensive and difficult to execute the necessary trades when they are needed most.

II.2 Credit Risk in Derivatives

The second broad category of financial risk exposure for those who trade derivatives is credit risk. Like market risk, credit risk is a familiar concept for fundamental securities, but its manifestation in derivatives is more complex. If a bank makes a loan of $1 million and the borrower declares

---

[3] See Barron's, Nov.3, 1997, p. MW 17.

bankruptcy the following day, the bank can lose the full $1 million. Similarly, if an investor buys $1 million of bonds or stocks and the issuer becomes insolvent, the full amount can be lost.

By contrast, if the bank had entered into a forward contract or a swap with a notional principal of $1 million and the counterparty became insolvent the next day, the loss might be negligible. Normally the bank could simply return to the market and replace the defaulted contract with another at approximately the same terms; the value of the underlying asset for a derivatives contract is never at risk. For example, it is not possible to lose $1 million, or anything close to that amount, from a default on a $1 million notional principal swap contract.

Depending on the particular instrument involved, there need be no loss at all from the failure of a derivatives counterparty, unless the market has moved away from the initial level, and in the direction that favors the solvent counterparty. Both forwards and swaps are normally set up to have zero value at the outset, so there is no credit risk at first. However, over time market prices relevant to the contract will tend to drift away from their initial values and one side of the contract will develop an embedded profit, while the other side will have a matching embedded loss. If the holder of the unprofitable position defaults at that point, there is a loss to the counterparty equal to the cost of replacing the defaulted contract at its original terms in the current market environment. Thus, a derivatives contract may involve little credit risk at the outset, but entail an exposure that grows over time.

How credit risk impacts derivative instruments varies across different types of contracts. A futures contract establishes a binding commitment to trade a fixed quantity of the underlying on the maturity date, at a price that is set today. When the trade is cleared, the original bilateral contract is broken into two contracts, each between one of the counterparties and the exchange Clearing House. The Clearing House is a AAA rated credit, so risk of default is minimized. Moreover, all contracts are collateralized by initial margin deposits, that are adjusted daily through the mark to market process to maintain adequate coverage of the potential credit risk exposure over the immediate future. Thus, there is effectively no credit risk on futures positions.

Forward contracts are like futures in that the contract specifies a commitment to a future transaction that is binding on both counterparties. The price is set so that the contract has zero value at the outset. However, as over the counter instruments, forwards are not automatically collateralized and marked to market, so credit risk exposure can build up over time. The expected future credit risk on a given trade depends on several factors, of which the volatility of the underlying asset is one of the most important. To see this, let us look at a forward contract tied to LIBOR.

We will first consider the case of a 3 year forward rate agreement on 6-month LIBOR with $1 million notional principal. The current level of LIBOR for all maturities is 6.00 percent and the forward is struck at the money. Thus, the counterparty with the long position enters into a 3 year commitment. At maturity, the strike rate of 6.00 will be subtracted from the actual value of 6-

month LIBOR in the market. If LIBOR has gone up, the short pays the long a cash amount equal to the difference between 6 months interest on $1 million calculated at the actual LIBOR and at 6.00 percent. If the rate has gone down, this calculation will yield a negative amount, meaning that the long pays the short.

For example, suppose the rate ends up at 7.00 percent. The short will pay the long (0.0700-0.0600) x 1/2 x $1,000,000 = $5,000. If the rate were 5.00 percent, the long would pay $5,000 to the short. Thus, if the rate is 7.00 at maturity, the long has a $5,000 credit exposure with respect to the short, but if the rate is 5.00, there is no credit exposure. (The short now is at risk that the long will default.) Obviously, under any reasonable interest rate scenario, neither counterparty can lose the $1 million notional principal or any amount near that.

Prior to maturity, the replacement value of a $1 million forward rate agreement struck at 6.00 percent will depend on the current market level of LIBOR and the possible rates that might occur at expiration. If the contract has positive value for the long at the current rate, there will be a credit exposure, while if the value is negative, the long has an embedded loss, no credit risk exposure, and would actually be made better off if the counterparty were to fail in such a way as to remove the long's liability.[4]

The time pattern of expected risk exposure for a forward with no initial market value starts at zero and rises over time as maturity approaches. The expected credit exposure at each future date is calculated as the expected value of the contract's market value on that date over all levels of LIBOR for which the market value is positive. Since the expectation is computed over only one side of the probability distribution for future interest rates, higher volatility that leads to wider dispersion in rates will produce greater expected credit risk.

We have built a simple binomial interest rate model to simulate the possible paths of future interest rates and the values of several types of interest rate derivative instruments over a 3-year horizon under the conditions described above. The model has one time step per month, and over a single step, the rate can go to only 2 possible levels, up or down by one standard deviation. We assume the annual volatility of the rate is 20%, making the volatility per time step equal to 20 x $(1/12)^{1/2}$ = 5.77%. This is a reasonable value for LIBOR volatility, that gives steps of about ± 35 basis points starting from an initial level of 6.00 percent. As is done in the binomial framework, we start by constructing a lattice of interest rates for the 36 months out to expiration. At expiration the derivative's payoff is computed for every interest rate in the tree on that date. The valuation then is rolled back through the lattice, to give the theoretical value for each date and rate all the way back to the present. Once we have this derivative value tree, the credit exposure

---

[4] As we will discuss further below, there remain legal questions in some jurisdictions about whether a counterparty's failure may remove the solvent counterparty's liability to make payments on derivatives transactions that have gone in the insolvent counterparty's favor.

for a given date is computed by multiplying the forward's value at each node for which the value is positive by the binomial probability of reaching that node, and adding them up.

Figure 3 plots the resulting estimates. The lowest curve shows the time pattern of risk exposure for the forward contract we have just described. For a $1 million notional principal contract under these assumptions, expected credit risk begins at 0 and rises over time to about $4,200. The slight waviness in the line is spurious. It is a result of approximation error in the binomial model we are using--a kind of model risk. Note that this calculation gives the mean value of the credit exposure, not the worst case.

The other common kind of derivative with a single maturity date is an option. A call option on LIBOR with a strike level set at 6.00 percent will pay the long the same amount as the forward if the market rate at expiration is above 6.00 percent, and zero for rates below 6.00. Because the long can receive payments at maturity if the rate moves up but does not have to make any payment to the counterparty when the rate has moved against her, the option has a positive value at the outset. The long must buy it from the option writer at the beginning, and in return, the writer must stand ready to make the payments called for at maturity if the option ends up in the money.

Since the option has positive market value at time 0, unlike a forward, the long has a credit risk exposure with respect to the writer from the outset, while the writer has all of the future liability under the contract and never has any credit risk exposure with respect to the option buyer. The dashed line in Figure 3 shows the time pattern of expected credit risk exposure for the option buyer. Under the assumed parameter values, an immediate failure of the option writer will produce a loss of $3,340. The expected value of the credit exposure rises gradually to $4,200 by expiration date. This is the same as the credit exposure for the forward because the option pays off for interest rates above 6.00 percent, which are the same rates at which the long to the forward would be owed a payment from the short.

An interest rate swap specifies a fixed "swap rate" at the outset, corresponding to the strike rate for a forward, and a set of future payment dates. On each date, the difference between the floating rate (the current market level of 6-month LIBOR in this case) and the swap rate is computed and the difference is applied to the notional principal. If the floating rate is above the swap rate, the "floating rate payer" pays the calculated amount to the "fixed rate payer," and if the floating rate is below the fixed swap rate, the cash flow goes the other way. It should be clear that a swap is simply a sequence of forward rate agreements with the same strike rate. The fixed rate payer is the long in the forward contracts and the floating rate payer is the short.

We have computed the pattern of credit risk exposure for a $1 million 3-year swap with payment dates every six months. The final swap payment occurring in 36 months is exactly the same as the forward we looked at before, and earlier payments correspond to five forwards with maturities of 6, 12, 18, 24, and 30 months. The dash-dotted line in Figure 3 shows the results. Since there are now six payments rather than one, the total credit exposure for the swap must be greater than for the forward until all but the final one remains outstanding. But since the swap rate is set to give zero value at the outset, there is no credit exposure at the beginning. There are then two factors operating in opposite directions as time elapses. Like a single forward contract, each outstanding component forward of the swap increases its exposure over time as rates drift away from the initial level. However, every six months one of them reaches maturity and pays off, which reduces the remaining aggregate risk exposure. This produces a humped total exposure curve, with a sawtooth pattern that is created by the periodic payments. The expected maximum exposure in this case is $11,200 and occurs in month 12 (just before the second payment is made). This curve shows greater choppiness than it should, again due to the effect of approximation error in the model.

Finally, we consider an interest rate cap contract. Similar to a swap contract which is a sequence of forwards, a cap contract is a sequence of call options with the same strike interest rate. Each individual option, known as a "caplet," entails a payment from the writer to the option buyer if 6-month LIBOR is above the strike rate on the periodic payment date. Like a single call, the cap contract requires a cash payment from the buyer to the writer at the outset, and subsequently only the buyer has credit risk exposure. The time pattern, shown in Figure 3 as the dotted line, is greatest at the initial date and falls in stair steps as the individual caplets mature.

There are several important points that should be clear from this discussion. First, credit risk in a derivative is much smaller than for an outright position in the same notional principal amount of the underlying asset. A $1 million loan for three years would have expected credit exposure of $1 million at every date, while the largest exposure among these derivative contracts was under $16,000 for the cap at time 0. Second, the time pattern of exposure is complex, it is significantly different for different instruments, and it depends on the volatility of the underlying in an important way, and frequently on a variety of other market factors. In fact, we have simplified this example considerably from what would have to be done to assess credit risk exposure for actual LIBOR-based derivatives. For example, by assuming LIBOR at all maturities was 6.00 percent at the outset, we eliminated a set of computations that would normally be required to take proper account of the current term structure and its future evolution. Finally, once again, it is apparent that this kind of analysis requires heavy use of mathematical models.

II.3 Operational Risk in Derivatives

Operational risk refers to the risk of a loss to a firm that is related to the internal handling of transactions and positions: the accounting and other back office operations supporting the firm's derivatives activities. Some of the more spectacular derivatives-related losses of recent years were primarily the result of operational risk.

A prime example was the huge losses mostly on futures and options tied to the Japanese stock market, that led to the collapse of Barings Bank, PLC. in 1995. The trades had been entered into by a 28 year old trader in the bank's Singapore office, who had been given overall responsibility both for trading derivatives and for the back office support of the trading desk. After losing a substantial amount for the firm's proprietary account through speculative trading in the Japanese derivatives market, he was able to conceal the losses in a bogus "customer" account while he increased the size of the bets he was making with the firm's capital in an attempt to recover the funds. The larger trades also went bad, causing increasing losses. By the time these were uncovered, they amounted to more than the total net capital of the bank, which was bankrupted.

Although operational risk can affect all of the trading operations of a financial firm, there are two factors that may make derivatives more susceptible to problems of operational risk than fundamental securities are. First, derivatives are generally much more leveraged. Since the object of contracts like futures, options, swaps, and similar instruments is to fix the terms for a transaction that will only take place in the future, a trader can make very large commitments on behalf of the firm and expose it to very great risk, with only a small amount of cash being exchanged in the present. For "cash market" transactions in fundamental securities, it is not easy to take large risk positions without there also being large concurrent cash flows that tend to make the situation visible to others in the firm. Although it is hard for a firm to protect itself completely against a determined fraud, uncontrolled trading that produces large losses is much harder to conceal in the cash market than in derivatives.

Second, since derivatives are more complex and less likely to be fully understood by non-specialists, senior management may be more disposed to leave the back office work and risk management of the derivatives business to the traders who seem to understand what they are doing. As the case of Barings shows, this is an extremely dangerous practice.

The recommendations from the Group of Thirty, as well as the other studies, with regard to dealing with operational risk in derivatives trading are clear and comprehensive. The two major themes are that senior management must be fully involved in setting risk standards and in overseeing the risk management system in their firms, and that there must be a clear separation of authority between the trading desks and the back office support and risk management functions. Implementing these recommendations throughout the financial services industry would eliminate much of the operational risk in derivatives trading.

II.4 Legal Risk in Derivatives

Legal risk is a problem whose essence is lack of information: Counterparties enter into a contractual agreement without full understanding of its legal aspects, and at a later date it turns out that the terms of the contract are not legally enforceable. This may come as a surprise to both counterparties (though only one will be damaged, while the other is relieved to be let out of a losing trade). The misunderstanding might simply be due to ignorance of a possibly obscure point of law applying to a particular transaction or type of transaction. More commonly, it arises because it has not been firmly established how the existing legal framework applies to a new type of instrument or trading strategy.

Legal risk is potentially present in all kinds of securities trading. Once again, however, it impacts derivatives particularly. We will not attempt to cover the legal risk faced by derivatives dealers in detail here. However, it is worth pointing out a few of the special ways in which it affects them.

One of the largest losses to dealers in derivatives in the 1990s occurred on a number of swap contracts they had entered into with several municipal governments near London. The largest of these were the boroughs of Hammersmith and Fulham. When rates went against the boroughs, the British courts, culminating in a 1991 decision by the House of Lords, ruled that they had exceeded their legal authority in entering into such contracts, and the swaps were void. This left the dealers to absorb losses of about $178 million on the defaulted contracts. Legal risk from lack of understanding of this law produced major losses for derivatives dealers. It is unlikely that financial transactions involving more traditional kinds of securities would have been as susceptible to this legal risk.

Examples of legal risk stemming from the lack of a fully established legal framework for new kinds of instruments are not uncommon. The U.S. regulatory system for derivatives divides responsibility among several authorities according to the type of instrument involved. Futures are regulated by the Commodity Futures Trading Commission (CFTC), options are regulated by the Securities and Exchange Commission (SEC) and the Federal Reserve and the Comptroller of the currency regulate banks and many of the activities of government bond dealers. This presents problems as new derivative contracts are created with elements that span more than one class of instrument. The CFTC found it necessary to resolve a significant source of legal uncertainty regarding whether swaps, traded extensively in the over the counter market, were actually a kind of futures and therefore requiring oversight by the CFTC. The CFTC's legal right to exempt swaps from oversight was made explicit by act of Congress only in 1992, and the CFTC did so in 1993.

Another source of legal uncertainty in the handling of derivatives transactions internationally concerns the enforceability of bilateral netting agreements among dealers. The general principle, which is well established in the U.S. and most other countries where major derivatives business is currently being done, is that when two firms are counterparties in multiple derivatives contracts, if one should become insolvent, the values of all of their outstanding contracts with each other are netted together. For example, suppose Firms A and B have two outstanding swaps with each other, one of which has market value of $300 million in B's favor and the other is worth $400 million in A's favor. If B becomes insolvent, under bilateral netting the two swap values will be combined, and B will owe A $100 million. In some countries, however, it has not been established clearly that this kind of agreement is enforceable in practice. One alternative would be that B, as the insolvent party, is protected from having to pay the $400 million to A, while A is still obliged to pay B on the $300 million swap. Uncertainty of this sort plays havoc with credit risk calculations. One of the specific recommendations of the Group of Thirty report on derivatives was a call for all countries to guarantee the enforceability of bilateral netting agreements.

A new source of legal risk has also been developing recently in the U.S., in response to several large losses involving derivatives. The risk stems from the possibility that a transaction which would seem to fall within an established set of legal principles may be called into question in a lawsuit. When a case involving complex derivatives is argued before a jury and a judge, who can not be expected to have a complete understanding of the financial instruments, the outcome is uncertain. There is a significant risk that an incorrect decision may be reached in such a case leading to a loss--perhaps unjustified, but a loss nevertheless. There is also the serious problem that for a derivatives dealer simply to have to defend itself against an extended public attack on its business practices and behavior can be very costly in terms of legal fees, a major distraction to the firm's personnel and very damaging to its client relationships. A firm facing such a lawsuit may well feel that it is completely in the right on the legal questions, but that winning in court is not worth the cost. The threat of a lawsuit by a disgruntled counterparty may be sufficient to make the firm accept a loss in an out of court settlement.

Legal risk from the threat of lawsuits presents a difficult problem that applies especially, though not exclusively, to derivatives. The recent dispute between Procter & Gamble (P&G) and Bankers Trust (BT) presents an example of how such a case can be made. Procter & Gamble entered into a nonstandard kind of swap agreement whose terms were fairly complex.[5] The nature of it was that if interest rates stayed low P&G would obtain financing at a very attractive rate. But with sharply rising rates, they were committed to pay interest at a floating rate that could quickly increase to extremely high levels, well over 20 percent annually. In fact, during the winter of 1994, rates did go sharply higher and P&G ended up realizing a $157 million loss on the position.

P&G claimed that they had not fully understood the terms of the contract and that Bankers Trust, not P&G, should be liable for the losses they incurred, because BT had not revealed relevant information. During the ensuing few years, the case was widely discussed publicly, damaging information was revealed, and BT found itself in difficulty with its public image and it customers, not to mention bearing mounting legal costs and loss of valuable personnel. In the summer of 1997, BT settled out of court, paying P&G about $150 million to end the case.

Without attempting to pass judgment on the rights and wrongs of this specific case, the known facts illustrate a new legal risk applying particularly to derivatives dealers. Derivatives are inherently more complex instruments than fundamental securities. If a firm enters into a

---

[5] See Smith [1997] for a fuller description of the actual contract and the possible reasons for it.

derivatives transaction with a dealer firm and makes a large profit, there will be no complaints. But if it makes a loss, it has the option of attempting to gain some or all of the losses back from the counterparty dealer by a lawsuit. If the case goes to trial, the plaintiff can be expected to argue that the transaction was complex and confusing, and that the dealer did not fully reveal its riskiness. In such an argument, it does not hurt the plaintiff if the jury and judge are confused about the transaction and about derivatives in general, since they may then be sympathetic to the claim that the deal was not well understood. The defendant firm, on the other hand, must attempt to make the jury understand how the transaction worked and to show that the plaintiff could and did understand what they were doing. Of the two, it is obviously easier to confuse a group of laymen about derivatives than to explain a complex derivatives strategy to them.

The option to sue whenever a customer experiences a large loss introduces an important new kind of risk for derivatives dealers. It is potentially extremely damaging to the derivatives market, if its use becomes widespread. Dealers are in the weaker strategic position in such cases for two reasons: They have the harder case to make in court, and it is extremely costly for them to be involved in the suit at all. But the practice of settling quietly out of court can add huge costs to being in the derivatives business, that will be reflected in higher fees and spreads, and greatly restricted markets, as firms will be obliged to reject many potential customers. Yet because a given dealer facing a bad case will still frequently find it cheaper overall to settle than to fight it out in court, there is a serious "public goods" problem with regard to this kind of legal risk. Dealers as a group would be better off if a few of them went to court and established legal precedents about how much responsibility they bear for their customers' investment decisions. But individually, the economics of the situation will frequently dictate settling rather than fighting.

II.5 Derivatives and Systemic Risk

The kinds of risks we have been discussing are all experienced at the individual firm level. A failure to oversee the back office operations properly can cause major losses for a single firm, but the problem is not expected to extend very far into the financial markets in general. By contrast, one of the frequently voiced concerns about derivatives is that they could lead to "systemic" risk, risk that a problem originating in the derivatives markets might end up damaging the entire financial system. While avoiding unnecessary losses at individual firms is clearly desirable, it is much more important to avoid the possibility of system-wide damage, from any source. Concern about systemic risk may be used to justify a call for increased governmental regulation of the markets, or restrictions of various kinds on participation by certain classes of investors.

Clear evidence that derivatives trading contributes to systemic risk, or even a clear description of how a derivatives-related event could cause a systemwide problem has been lacking, consistent with Volcker's comment that "...systemic risks are not appreciably aggravated" by derivatives. Yet the importance of the issue keeps concern alive that derivatives activities may cause systemic risk in some unforeseen way either today or in the future.

This is a major issue that can not be fully addressed here. However, I will offer a perspective on the question to suggest that, by their nature, derivatives activities are less likely to cause systemic problems than is trading in fundamental securities. This is not meant to be a rigorous proof, but rather an intuitive discussion of the consequences of an important difference between derivatives losses and fundamental losses.

Derivatives are inherently different from fundamental securities with respect to systemic risk because they are a zero sum game. To illustrate why that can be expected to make a difference with respect to systemic risk, let us take an analogy. Consider a family that holds a nightly poker game around the kitchen table. From the perpective of an individual family member, the effects of losing $100 in poker and losing $100 when his wallet is stolen in the subway are similar: his wealth falls by $100. But from the perpective of the "system," in this case the household as a whole, there is a lot of difference between these events. The loss in the poker game simply transfers $100 from one family member to another; it has no impact on overall household wealth. But losing $100 in the subway is a net reduction in total household assets. The latter can be expected to have more serious potential consequences for overall household welfare.

As far as the financial system is concerned, derivatives losses are like the losses in the family poker game. A dollar lost by one market participant always produces a dollar gained by someone else. This means derivatives risk can only be redistributive risk--money changes hands but there is no loss of wealth to the system as a whole. By contrast, events such as an earthquake that destroys a town, a firm's investment project that fails, or a decline in a nation's overall productivity that drives the stock market down, all produce losses in aggregate wealth to the whole economy. These are the kinds of risks that fundamental securities are exposed to. Losses of this kind may be shared through diversification, they may be transferred by insurance contracts (or by hedging with derivative instruments) from those less able to bear the risk to those more willing and able to do so, but however it is redistributed, the entire loss must be borne within the system.

Therefore, I suggest that by the nature of a derivative contract as a zero sum game, derivatives risks inherently entail less chance of systemwide disruption than do the risks on fundamental securities that correspond to losses of wealth to the whole system.

III. Model Risk in Trading Derivatives

One of the most apparent differences between trading in derivatives and trading in fundamental securities is the enormous importance of theoretical models and sophisticated mathematical tools to the derivatives market. Derivative instruments themselves tend to be more complex than fundamental securities, and what are now standard operating procedures for trading them and managing the associated risks involve much greater complexity.

An important reason for this is that valuation theory is fundamentally different for the two kinds of financial instruments. Theoretical valuation models for fundamental assets are deduced from considerations of equilibrium in the capital markets. The principles at work are well-established (e.g., investors like high returns and dislike risk) but they do not produce very precise valuations. For example, a risky security should have a higher expected return than a risk free security. But how much higher? Pricing theory can not give a specific value; it depends on the degree of risk aversion in the investor population. This dependence on unobservable expectations and risk aversion makes precise valuation difficult for fundamental securities. It also means that market participants do not rely very much on theoretical models in trading them.

Pricing models for derivatives, on the other hand, value a derivative instrument relative to its underlying asset. The underlying need not be priced "right" in the market for us to say with considerable confidence how a derivative instrument based on that asset ought to be valued given the price of the underlying. This is possible because the models are derived from considerations of arbitrage between the derivative and its underlying. If the derivative's market price differs very much from the model value, there will be an arbitrage strategy that earns riskless excess returns, regardless of how the underlying is priced vis-a-vis other securities in the market.

The pricing model for forward and futures contracts is quite straightforward. It is based on the cost of buying the underlying asset, hedging its risk by selling it forward, carrying the hedged position through time until futures maturity, and then delivering. The resulting "cost of carry" pricing model is applicable for all forwards and futures for which the underlying is storable, i.e., it can be bought in the present and held over time. Financial instruments, including long term government bonds, equities and foreign currencies are good examples of storable underlying assets; live animals, commodities like oil and electricity for which supply comes from current production, and economic variables, like the Consumer Price Index, that are not prices of traded securities are not easily storable, and are not governed by the cost of carry relation. For markets in which the model applies, the possibility of arbitrage between the derivative and its underlying enforces theoretical pricing in the marketplace, at least to an approximation, whose degree of closeness will be a function of how hard the arbitrage trade is to execute.

Theoretical models for pricing derivatives with option features are significantly more complicated. The celebrated Black-Scholes [1973] formula is also derived from an arbitrage-based trading strategy that combines the option and its underlying to create a delta neutral portfolio that is free of market risk for small changes in the underlying price over a short time interval. As we discussed above, however, in practice a delta neutral hedge of an option against its underlying is far from completely risk free, and keeping it delta neutral over time requires what can be a very active rebalancing strategy. In fact, while market pricing of options depends heavily on theoretical models derived from such dynamic hedging strategies, there is virtually no trading in actual markets that attempts to execute exactly the continuous rebalancing option replication strategy that these models are based on. For one thing, the theoretical strategy involves an enormous number of transactions, but trading costs in practice are too large for that to be feasible.

Nor is continuous trading possible even if one wanted it, when markets are closed at night and on weekends. Nor are potential arbitrageurs certain that they know the exact values of volatility and other model parameters.

Option pricing formulas require assumptions about the stochastic process that the price of the underlying asset follows. The form of the distribution must be specified and the numerical values of some of its parameters must be known. The most common assumptions are that the asset price follows a lognormal probability distribution with a known volatility.[6] The Black-Scholes model and the variants derivable from it embody these assumptions. It is widely used in the marketplace, both for pricing options and for estimating the Greek letter risk exposures. Yet even while the model is used actively, traders are also well aware that its assumptions do not hold in practice. Actual security returns are not really lognormal and volatility is not known, but must be estimated subject to considerable error.

In this section, we will consider the risks in trading derivatives and managing derivatives risks that result from the need to use theoretical valuation models. We will focus on several classes of model risk. The first is risk that comes from using wrong models, in particular assuming lognormality when the true distribution has "fat tails." The second type of risk arises even when one has the right model, because necessary input parameters such as volatility can only be obtained as forecasted values, subject to error. Third, beyond risk from the use of models that do not exactly correspond to reality, there is implementation risk, that programming errors and other extra-model problems will cause incorrect answers even though the model is essentially correct. Finally, there is the risk that a model may give correct theoretical values, but the market prices will be different.

III.1 Wrong Models

There are many ways a valuation model can go wrong. Indeed, derivatives practitioners know well that their models are not exactly correct, and will try to make accommodations for model problems that they are familiar with. To deal with the fact that the Black-Scholes model assumes the underlying asset's volatility is a known constant, while in the real world, the true value of volatility is quite uncertain, traders will frequently use "implied" volatility obtained from market option prices as the model input. Implied volatility (IV) for a given option is obtained by solving a valuation model backwards to find the volatility that would make the model value equal to the observed market price.

---

[6] A random variable follows a lognormal distribution if its logarithm has a normal distribution. This is (almost) the same as assuming that the asset's rate of return over short intervals is normally distributed.

Of course, this procedure is circular: one assumes the market price is the true value, and then fiddles the volatility input parameter until the model produces that value. Worse still, the procedure does not work exactly right. The underlying is one asset and can have only one actual volatility. But, there are typically numerous options with different strike prices trading on the same underlying asset, and each of them will produce a different implied volatility. The existence of multiple implied volatilities shows that option prices in the market are not consistent with the valuation model.

The array of implied volatilities for different strike prices has a customary pattern known as the volatility "smile." When IVs are graphed against option "moneyness" (how far an option is in or out of the money) options that are at the money normally have the lowest implied volatilities, while those on the wings have higher values, producing a generally smile-shaped pattern. Although a smile means option pricing in the market is inconsistent with the model, traders still tend to use a familiar though incorrect model and try to adjust in various ways for its known shortcomings. Such adjustments are simply rules of thumb, however; they are not rigorously justifiable.

This leads to model risk, the amount of which is not easily known. To the extent that risk managers, being further removed from the market than traders, are less able to make the subjective corrections necessary to make the models work, the importance of model risk may not be as well understood as the other types of risk described in the previous section. Theoretical option valuation models do not fully capture how the market actually prices options, but the reasons for the discrepancies are not directly observable. We will now discuss several of the most important important sources of model error.

"Fat Tails" in the Distribution of Asset Returns

The lognormal probability distribution is widely used in option pricing models as the assumed stochastic process for the price movements on the underlying asset. There are theoretical and practical reasons for this. On theoretical grounds, the lognormal arises naturally when asset returns follow a random walk with a constant variance as the interval between price observations goes to zero .

Equation (1) describes the returns process for an asset price that follows a random walk.

$$\frac{S_{t+1} - S_t}{S_t} = r\,\Delta t + \epsilon_t\,; \qquad E[\,\epsilon_t\,] = 0 \qquad VAR[\,\epsilon_t\,] = \sigma^2\,\Delta t \tag{1}$$

In period t, the asset's expected return is given by an annualized rate r multiplied by the length of the time interval Δt. The actual return is the expected return plus a zero mean random disturbance ϵ_t whose variance, $\sigma^2 \Delta t$, is proportional to the length of the interval. The random disturbance is independent across time periods.

The limiting process, the lognormal diffusion that is the foundation of the Black-Scholes option pricing formula and many subsequent variants, is given by

$$\frac{dS}{S} = r\,dt + \sigma\,dz\,; \qquad \begin{array}{l} E[\,dz\,] = 0 \\ VAR[\,dz\,] = 1\,dt \end{array} \tag{2}$$

The random component of returns denoted dz still has zero mean and a variance of $\sigma^2 \Delta t$ over a finite time period of length Δt, and it is independent over time. One difference is that in going to the limit, the disturbance converges to a normal distribution, which induces a lognormal distribution for the asset price and for its returns over finite time periods.

The random walk model has received a great deal of support from empirical researchers over time, so it is natural to extend it to the continuous-time world of option pricing. On practical grounds, the mathematics of lognormal distributions is well-known and relatively simple, at least in comparison with the alternatives. Option pricing models that drop this assumption tend to become significantly more complex. Greater complexity typically allows a model to achieve a better fit to the available data (a smile pattern in implied volatilities can be partially flattened out, for example), but at the cost of greater fragility in forecasting (the model's performance tends to degrade more rapidly when it is used out of sample). Moreover, non-rocket scientist model users often have some familiarity with normal distributions, which extends to the lognormal, but very little intuition about more complex distributions.

Yet, one feature of actual returns distributions that is observed in all sorts of financial markets is "fat tails." Empirical returns distributions exhibit more very large values, both positive and negative, than a lognormal distribution with the same mean and variance. That is, they have more probability weight in the tails of the distribution. In order for there to be more very large deviations from the mean without an increase in the variance, these fat-tailed distributions also have more very small returns than the lognormal, and fewer middle-sized ones. We will be considering percentage returns over short intervals which will have a normal distribution when the prices are lognormal.[7] Figure 4 illustrates the difference between a normal and a fat-tailed

---

[7] Strictly speaking, for a diffusion model like equation (2), the continuously compounded return has a normal distribution while the percentage returns over a discrete interval will be lognormal. For returns over an interval as short as one day, there is no appreciable difference between the normal and the lognormal distribution.

distribution.

Fat tails can arise from a variety of causes. One is the existence of "jumps" in prices. In a diffusion model like equation (2), all randomness is contained in the dz term so the variance of price changes goes to zero as the observation interval shrinks. This means the price process becomes continuous, ruling out the possibility that it could jump from one level to another without passing through every value in between. Yet actual financial series like exchange rates and interest rates may do just that, which produces larger moves over short time intervals than the lognormal allows for.

Beyond inducing a nonlognormal returns distribution, price jumps cause more fundamental problems for derivatives pricing models. Because it is impossible to rebalance a delta neutral hedged position continuously when the price jumps, the ability to replicate the derivative's payoff exactly with a dynamic trading strategy is eliminated. This undermines the whole principle upon which the pricing formula is based. Alternative option pricing models that allow price jumps have been devised, but they are less satisfying than models based on pure diffusion processes: the jump risk can not be hedged, so option values must incorporate an unobservable market-determined premium for jump risk. This eliminates one of the great strengths of the Black-Scholes model, that theoretical option values do not depend on unobservable preferences. Practitioners typically use pricing models that do not include jump risk, and try to adjust the model values subjectively in cases where it is expected to be especially important, such as short maturity out-of-the-money options.

Fat tails in the returns distribution may come from other sources as well. Equation (2) assumes the volatility at each instant is a constant (annualized) value σ, but the empirical evidence shows that realized volatility for actual financial variables varies over a broad range. Uncertainty about the future volatility makes derivatives pricing and risk management into a forecasting problem. In the next subsection we will examine in detail how volatility forecast errors translate into risk for derivatives traders. Here we simply observe that time-varying volatility induces nonlognormal returns distributions, and invalidates the use of a simple BS-type model.

Various lines of attack have been followed in attempting to extend the standard framework to stochastic volatility. One simple approach is to focus on the empirical observation that volatility appears to be related to the level of the asset price, going up when the price falls and down when it rises. This can be accommodated with minor modification to equation (2) by allowing the volatility parameter σ to become a function of the asset price, $\sigma(S)$. The volatility is then random but the model still has only one source of risk, dz. More elaborate models of this kind employ the GARCH (Generalized Autoregressive Heteroskedasticity) framework to model volatility changes.

A more general class of stochastic volatility models allows volatility to evolve independently from asset prices by introducing a second stochastic variable. Unfortunately, like allowing for jumps, this creates an unhedgeable risk in the option price and eliminates preference-free pricing. Option valuation models become quite difficult computationally under any sort of stochastic volatility. Moreover, introducing new unobservable parameters associated with the volatility process into the valuation model makes the estimation and forecasting problems of putting the model into practice much more severe. Again, traders are aware that volatility changes unpredictably over time, but they continue to use fixed-volatility pricing models with subjective adjustments.

These are known problems with the standard derivatives valuation paradigm. Others could easily be added to the list: stochastic variation over time in other input parameters like interest rates and dividend payouts, random fluctuations in the parameters of the volatility process itself, variation in the market prices of unhedgeable risks, and so on. Such things are to be expected: after all, a model is just that, a model. It is inherently oversimplified, representing a compromise between realism and tractability. It would be naive to expect great accuracy from a theoretical model of derivatives prices, or of any other economic relationship. However, the size of the model error and the proportion contributed by different sources is not widely known.

Use of theoretical models is becoming more prevalent outside trading rooms, for purposes for which it is not easy, nor necessarily appropriate, to apply the subjective adjustments that permit traders to correct for their significant shortcomings. Firmwide risk management will be determined by estimated exposures to different risk factors, as obtained from pricing models; Value-at-Risk computations will be made by putting estimated volatilities and correlations into lognormal distributions; risk exposures will be reported and regulatory risk capital standards will be set based on estimated probabilities of extreme events, as calculated from normal or lognormal distributions; and so on. In all of this, there is little formal recognition of the fact that these seemingly rigorous calculations actually are made by putting inaccurately estimated parameter values into incorrect theoretical models.

To illustrate the impact of fat tails in actual returns distributions on risk calculations using a lognormal distribution, we examine estimates of "tail events" for four major financial variables that might have been made using data available at the time. We obtained daily data for 3 month LIBOR, the U.S. Treasury 10 year bond yield, the Standard and Poor's stock index, and the dollar / yen exchange rate, for 750 trading days, ending Dec. 31, 1996.[8] Starting on day 251, we computed the sample value of σ from the previous 250 days' prices and used it to estimate the 5% and 1% tails of the returns distribution under the assumption of lognormality. An observation falling in the 5% tail is generally taken to be a rare event or outlier in statistical work; the 1% tail is both used in research as a more stringent cutoff point, and more importantly for this discussion, it has become widely adopted as an appropriate benchmark for measuring Value-at-Risk and for risk reporting by banks.

[8] Data were obtained from Datastream.

Treating a fat tailed distribution as if it were lognormal leads to two related but different types of estimation errors. First, setting the critical value for the x% tail based on a lognormal with the same variance as the actual distribution can not be expected to capture exactly x% of the actual observations. The true distribution will have more weight in the extreme tails than the lognormal, but less at intermediate values. Thus, there could be more or less than 5% of the actual observations that are below the mean by greater than the 5% critical value for the lognormal (1.64 standard deviations). For the more extreme 1% tail, the fat tail of the empirical distribution should become apparent, so we expect to find that more than 1% of the realized returns exceed the lognormal's 1% critical value of 2.33 standard deviations away from the mean.

The second difference between a fat tailed distribution and the lognormal is that those values that do end up in the tail will tend to be more extreme. The average across all observations in the tail will be greater for the empirical distribution than for the lognormal. For example, the average 1% event will be worse in reality than the lognormal would suggest. Table 2 shows both aspects of model error in using an estimated lognormal distribution in actual markets.

If S follows equation (2), its logarithm has a normal distribution. In derivatives pricing formulas, the actual probability distribution for the underlying is replaced by the "risk neutral" distribution that has the same volatility but a mean return equal to the riskless interest rate. The diffusion equation is just equation (2) with a redefinition of r. For an underlying whose total return comes from price appreciation, r is the risk free interest rate instead of the mean return on the underlying. If the underlying pays a dividend or some other cash distribution at a continuous rate q, r is the riskless rate minus q, while if the underlying is an interest rate, r is normally set to zero. For simplicity, we will set r to zero for all of the instruments we consider. The results will therefore only be approximate, but for the short time interval of one day the difference is negligible, because the results are dominated by the random component of returns.[9]

Assuming $\ln(S_1/S_0)$ has a normal distribution with mean 0 and annualized standard deviation σ, for the return over a time interval Δt, the cutoff for the lower x% tail of the distribution will be a number α of standard deviations below the mean, such that

$$N\left[\frac{\ln\frac{S_1}{S_0} + \frac{\sigma^2}{2}\Delta t}{\sigma\sqrt{\Delta t}} + \alpha\right] = x \tag{3}$$

where N[.] denotes the cumulative normal distribution. The upper x% tail is obtained by solving

[9] For example, if the true r is 15 percent and volatility is 0.1500, setting r to zero for calculations at the one-day interval would increase the measured volatility to 0.1503.

(3) for the α value that gives a lower tail probability of (1 - x).

We will also be interested in the size of the average tail event, that is, the expected value of the percentage change in S (i.e., $S_1 / S_0 - 1$) for all changes large enough to exceed the critical value. Under our assumptions, those values are given in terms of standard deviations by equations (4) and (5).

Define

$$d_1 = -\alpha + \frac{\sigma\sqrt{t}}{2} \qquad \textit{and} \qquad d_2 = -\alpha - \frac{\sigma\sqrt{t}}{2}$$

Then,

$$\textit{Expected value in the } \alpha\% \textit{ lower tail (in standard deviations)} = \frac{1 - N[\,d_1\,]}{1 - N[\,d_2\,]} \qquad (4)$$

$$\textit{Expected value in the } \alpha\% \textit{ upper tail (in standard deviations)} = \frac{N[\,d_1\,]}{N[\,d_2\,]} \qquad (5)$$

The upper portion of the table presents results for the 5% upper and lower tails, and the lower portion does the same for the 1% tails. To enhance comparability across markets and dates, all results are expressed in terms of the (estimated) volatility, rather than in percents or dollars. To convert these into percents, they should be multiplied by σ_d, the standard deviation of the percent return over one trading day, where $\sigma_d = \sigma / \sqrt{255}$.

The tail cutoff values for a lognormal distribution and the theoretical mean value for observations that fall in that tail are shown for each subsection, followed by four columns of results. The first column gives the fraction of actual returns that would have exceeded the lognormal tail cutoff value. Thus, for 3-month LIBOR the estimated 5% lower tail captured only 2.80% of the actual changes (14 out of 500), while the T-bond yield fell by more than the 5% cutoff value on 5.80% of the days in the sample. The second column gives the average of the actual returns for those observations that did fall in the theoretical 5% tail. For example, under the lognormal the average 5% lower tail observation should be -2.06 σ_d, but for the 14 days classified as "5% tail events" for 3-month LIBOR, the actual average was -3.42 σ_d.

The third column shows where the cutoff point would have had to be placed to capture the desired percentage of actual returns: For LIBOR, to cover 25 out of 500 observations, the lower tail cutoff would have had to be -1.15 σ_d. Finally, the last column gives the actual mean return for observations in the desired tail of the actual distribution. By setting the cutoff at -1.15 σ_d, there would have been 25 tail events for LIBOR, with a mean return of -2.50 σ_d .

The results in this table combine model error from two different sources. There is parameter forecasting error because the tail cutoff points are derived from volatility estimates computed from past data. There is also distributional error due to the fat tails of the actual distributions. Estimation error is presumably random, so it should cause overestimates of the tail cutoff point as often as underestimates. The distributional error, however, will cause mistakes biased toward one side: observed extreme values will be more extreme than the estimated lognormal distribution accounts for.

The table presents results in terms of numbers of standard deviations, since the estimated volatilities of yield or price changes are different each day. But to get an approximate idea of what these discrepancies might translate to in terms of one day risk exposures, the standard deviation figures can be multiplied by the average 1-day standard deviation estimates in basis points or percents, as shown in the Note. For example, consider the 1% lower tail for the S&P 500 index. Over the 500 days in the sample, the average of the estimated volatilities for the S&P index was 0.599% per day. Taking that value as σ_d , the 1% tail was estimated to cover losses of more than -1.40% over a day. The average among losses at least that large was expected to be -1.59%. However, by setting the tail cutoff at -1.40%, there were actually 2.60% of the days classified as tail events, and the average loss among them was -1.89%. The 1% tail of the actual returns distribution contained the 5 days with losses exceeding -1.97%, and on those days the average loss was -2.39%. Thus, the combination of statistical error in estimating volatility from past data and the fat tails of the actual distribution of stock returns produced a substantial underestimate of the loss on the S&P 500 stock index portfolio that would be experienced with 1% probability.

Comparing results across markets and tail definitions, there are several regularities that are important in assessing and managing risk for these instruments. First, while three out of four markets exhibited fewer tail events at the 5% level than predicted, for the extreme 1% level only one tail for one market was right; the other seven cases showed "too many" tail events, and at least twice as many as predicted for four of them. Moreover, for those returns that were classified as tail events under the lognormal, the average size of the event exceeded the predicted average in every case. LIBOR presents an interesting example: There were relatively fewer tail events than expected, except for the 1% upper tail, but those that did occur were much larger than the predicted average under the lognormal.

The tail cutoff points for the desired 5% and 1% critical values in the actual distribution were not too far from those in the lognormal at the 5% level, but were distinctly more extreme at the 1% level, and in every single case the mean return for observations in the desired tail of the true distribution was substantially greater than what was indicated by the lognormal.

III.2 Wrong Parameter Values

The previous subsection focused on errors in assessing risk exposure due to assuming an incorrect probability distribution for underlying asset returns. A closely related problem is that even a correct model requires the user to input a value for the volatility of the underlying from the present through expiration day, as well as other variables like future interest rates and dividend yields, all of which must be forecasted using current information. Estimation error causes model risk when forecasted values are used in place of true values in a pricing model. Model risk will produce mispricing of derivatives and also inaccurate hedging calculations.

There are a number of procedures for obtaining a volatility estimate, that fall into three broad classes. The most basic is simply to calculate the realized volatility in a sample of recent price data for the underlying and assume that the same value will apply over the future life of the derivative one is pricing. Variations on this method involve the choice of how much past data to include, periodicity (e.g., daily data versus monthly data), whether deviations are measured around the sample mean or around an imposed mean value such as zero, and whether to downweight old data.

In Figlewski [1997], I examined the impact of these variations on forecast accuracy for volatility of the 3 month T-bill rate, the 20 year T-bond yield, the S&P 500 index and the Deutschemark / dollar exchange rate. A capsule summary of the findings with regard to the choice of volatility estimator is as follows.

* Out of sample forecast errors of even the best method tend to be quite large.

* The forecast error for longer horizons tends to be lower than for short horizons.

* Although in practice it is common to estimate volatility from quite short historical samples, using a larger amount of past data (e.g., several times the length of the forecast horizon or more) generally gives considerably greater accuracy, except when forecasting over the very shortest horizons (e.g., less than 3 months).

* Estimating from daily data improves accuracy for short horizons (6 months or less), but for longer horizons, monthly data gives better results because it is not affected so much by transient noise in market prices.

* Since the statistical properties of the sample mean make it a very inaccurate estimate of the true mean, taking deviations around zero rather than around the sample mean typically increases forecast accuracy.

As an example of the second point, for each month from January 1952 through December 1990 estimating volatility from 5 years of past monthly data and using it to forecast volatility over the next two years gave a root mean squared error (RMSE) of 4.17% when the realized volatility averaged 14.25%; RMSE for 3 and 5 year forecasts, respectively, were 3.62% and 3.10%.

The issue of whether to weight each sample point equally or to downweight old data was not addressed in that study, but we present some evidence on that question below.

Since the reason to forecast volatility is that it is time-varying, there is a logical inconsistency in using a framework that assumes constant volatility to estimate it. Another class of volatility estimators based on past data try to model the volatility process in order to take the current state of the system into account in computing the volatility. The most common of these approaches uses the GARCH framework. Figlewski [1997] also examined GARCH estimators. They were found to be useful primarily for short term forecasting of stock returns volatility with daily data; for longer horizons and in different markets they did not work as well as historical volatility. And for the cases in which GARCH was the most accurate estimator, the errors were still very large. For example, RMSE in forecasting daily volatility of the S&P 500 index over a three month horizon was 5.37% relative to the average realized volatility of 13.29%.

The other major way to obtain a volatility estimate is directly from the market prices of traded options. The implied volatility derived from an option's market price is felt by many practitioners, and academics as well, to be the best estimate. That implied volatility is a very accurate estimate of true volatility is far from established by empirical research, however. One large study, by Canina and Figlewski [1993], found that IVs from the S&P 100 stock index options market (one of the most active in the U.S.) appeared to contain no information at all about future realized volatilities. Figlewski [1997] reviews a number of empirical studies of the forecast accuracy of implied volatilities. The typical study shows that IV does contain information about future volatility, and generally more than the particular variant of historical volatility tested, but that IV is biased and forecast errors are substantial.

As we discussed above, there is an inconsistency in using implied volatility as a forecast of the underlying asset's actual volatility, when each option traded on the underlying typically produces a different IV value. Yet, each IV is obtained directly from an option's market price, so it may be the most useful volatility input to the pricing model when one is trying to assess short term market risk for that particular option. This is true regardless of whether the IV is a good forecast of the actual volatility of the underlying. IV is a measure of how the market is currently pricing the option relative to the underlying, and that relationship may be fairly stable over short time intervals even if it only means the option's mispricing today is a pretty good indicator of tomorrow's mispricing. IV may produce considerable model error, however, if it is an input to a model that is being used to value long term derivatives or for a purpose like computing Value at Risk that depends on an accurate estimate of the true probability distribution. Lack of a broad range of prices for traded options with different strike prices and maturities also limits the availability of implied volatilities in many cases.

The nature of the estimation error in predicting volatility from historical data is illustrated in Tables 3a and 3b. We have daily and monthly data on the same markets we examined above: 3 month LIBOR, the 10 year T-bond yield, the S&P 500 stock index and the yen / dollar exchange rate. For each market, several different forecasting horizons are examined, and for each, we try three historical sample lengths and also a forecast using exponentially declining weights.

To explain the details of the procedure, let us focus on the single example of forecasting the volatility of 3 month LIBOR over a 24 month horizon using 24 months of historical data, shown in Table 3b. Our data sample begins in January 1971. Since we need up to 5 years of past data to compute some of the estimates and we want to be able to compare the performance of all of the methods over the same period, the first forecast date is January 1976. In this case, continuously compounded changes for the observed interest rates from the last 24 months are computed by taking the first differences of their logarithms. Rather than calculating the standard deviation around an inaccurate sample mean, we impose zero as the mean value for the change in rates. The variance is then just the average of the squared log changes. This is annualized by multiplying the monthly variance by 12, and taking the square root gives the volatility. This estimate is then used as the volatility forecast going forward from January 1976; the historical volatility over the past 24 months just calculated becomes the first forecast for both the 24 month horizon and the 60 month horizon.

Next we compute the realized volatility over the 24 month period from January 1976 to January 1978, also constraining the mean to zero. The difference between the forecast and the realized volatility becomes the forecast error for January 1976. The forecast date is then advanced one month and the procedure is repeated. This continues until December 1991, the last date for which our data sample allows calculation of the realized volatility over a 60 month forecast horizon. The series of volatility forecast errors are then squared and averaged, and the square root taken to produce the root mean squared forecast error of 13.3%. This is the measure of forecast accuracy for the strategy of using 24 months of past data to predict future LIBOR monthly volatility over a 24 month horizon.

We note that the forecast errors computed in this way will not be serially independent, since the volatilities calculated for consecutive dates will come from almost exactly the same data points. Lack of independence will not bias the estimated RMSEs, although it would cause inconsistency in an estimate of the standard error of the RMSE. We think of this procedure as a way to assess the impact of estimation error for a financial institution that writes at the money calls every month, basing pricing and hedging strategy on volatility values estimated from past data in a standard way. Lack of independence will produce serial correlation in the pricing errors, which will show up in the form of runs of losing and winning months. To give an idea of how this impacts performance, in later tables we report both the worst single month and the worst year for the strategy.

The same procedure is used to fill in RMSE values for all combinations of market, historical sample and forecast horizon. The sample period labeled "All available" uses all past data back to the beginning of the sample as of each date. Thus, unlike the other methods, the amount of historical data in this volatility estimate grows over time, from 5 years on the first date to 20 years or more on the last. The table shows that this is an effective strategy in three of the four markets we consider.

The "All available" estimate uses every past observation weighted equally. However, it is generally felt that recent observations are more meaningful than ones from the distant past. A relatively easy way to take account of this is to weight each data point in inverse proportion to its age. The exponentially declining weight forecast is computed as in equation (6), where $0 < w \leq 1$ is the weighting factor and r_t is the log price relative from date t.

In the results shown here, the weights were computed for each date by analyzing only historical data that would have been available at the time, with the optimality criterion being out of sample RMSE. For each date, this required dividing the available historical data into an estimation sample and a forecast sample. The procedure we have adopted is somewhat arbitrary, but represents one plausible way for such a computation to be done.

As an example, consider the 24 month forecast horizon using monthly observations. To allow multiple past forecast horizons, on date t we compute realized volatility over 12 overlapping 24-month periods, {t-25 to t-1, t-26 to t-2, ..., t-36 to t-12}. Given a trial value for w, for each of those periods we compute a volatility forecast using equation (6) with all available data prior to the beginning of that period. This produces 12 (unfortunately overlapping) forecasts and forecast errors from which the RMSE is calculated. We then search over values for w to obtain the value that would have minimized RMSE. That is the w used on date t to forecast volatility over t+1 to t+24. We then advance a month and repeat the process to obtain a weighted volatility forecast for the period t+1 to t+25.

The difficulty arises at the beginning of the sample. To the extent possible, we would like to use an estimation sample of at least five years and a forecast sample allowing multiple periods of length equal to the forecast horizon we are trying to optimize for, so that a meaningful RMSE can be computed. That is not feasible in every case. To give a larger historical sample for the first forecast dates, we set w=1.0 for the first year with the daily data and for the 24-month horizon, and for the first 4 years with the 60 month horizon. This still can leave only 3 years of sample data to compute the volatility for the first of the forecasts that go into the RMSE we are testing to find the optimal w. This problem only affects the beginning of the sample, but it is inherent in the use of an exponentially declining weight forecast when the weight must be obtained by examining past data.

The optimal decay factor must be determined by searching over possible values. Although there is no strict necessity for the weight to be no greater than 1.0, a w above 1.0 would have the effect of increasing the importance of a data point as it ages, which makes no intuitive sense. It is customary to constrain $w \leq 1.0$, and we have done so here, with the result that some of the optimal decay factors turn out to be no decay at all, making them identical to the "All available" figures.

Table 3a presents forecast results for daily volatility over horizons of 1, 3, and 12 months using daily historical data. Table 3b does the same thing for monthly observations to forecast over 2 and 5 year horizons. In the daily table, for convenience in estimation a "month" is defined to be 21 trading days in all cases. Notice that the samples span slightly different periods. In particular, it is only necessary to hold out 12 months of daily data at the end of the sample for post-sample forecasting, rather than 5 years as in the monthly table.

One striking result is how large the forecasting errors are. For example, on average the forecast of 3-month LIBOR volatility over the next 24 months based on the last 24 months had an RMSE of 13.3 percent, which is very large relative to a mean realized volatility of 25.4 percent. Roughly speaking, this means that about a third of the time, the predicted volatility would be more than 50% above or below the true value. LIBOR is actually the worst case, but errors are substantial for all of these markets.

The results shown in these tables are consistent with those found in Figlewski [1997] for different sample periods and different markets. For monthly data, it is generally the case that the best estimates come from using the largest possible historical sample, but for daily observations and shorter forecasting horizons, performance is better if the historical sample is several times as long as the horizon, but not too long. Exponentially declining weights seem to increase forecast accuracy with daily data in some cases. Surprisingly, forecasting accurately over longer horizons seems to be easier than over shorter ones.

One interesting feature of these results is that (annualized) volatility for daily data appears to be substantially lower than for monthly data. This may be due partly to the somewhat different sample periods. It also may be a result of short term positive serial correlation in the data series. If price changes are not independent over time, estimated volatility will be affected. Positive autocorrelation, which occurs when observed prices adjust to new information with a lag over short intervals, reduces estimated volatility. This problem largely disappears with longer differencing intervals, which is the reason to use monthly rather than daily data for longer horizon forecasting.

In Tables 4 and 5, we examine the impact of volatility forecasting errors on a bank or financial institution that writes at the money calls each period (either every month or every day). The option pricing models we employ are all standard variants of the Black-Scholes model that are commonly used for these securities. The relevant equations are given in the Appendix. In each case, enough options are sold to produce $100 of premium income, so the performance figures shown in the table may be interpreted as percentages of the initial option price. Results are presented for all daily and monthly horizons examined in Tables 3a and 3b.

Table 4 looks at the strategy of selling the options at their model values, investing the proceeds, and simply holding the short position until maturity without hedging. The left side of the table uses the forecast based on historical data that had the lowest RMSE in Tables 3a and 3b as the volatility input.[10] The results therefore reflect model error due both to inaccurate volatility inputs and to shortcomings of the model. The right side of the table shows the results if the options are priced using the realized volatility over the option's life as the input to the model. This removes the effect of volatility estimation error, leaving only the impact of inaccuracy in the valuation model itself.

For each volatility, the first two columns give the mean return and standard deviation of the strategy. Since there is no hedging, the mean return on option writing should be a function of the expected value of the change in the underlying asset. For the stock market, this should be positive and well above the risk free interest rate. (Over the long run, stocks have averaged returns between 8 and 9 percent above Treasury bills, and as leveraged instruments, call options have higher expected returns than the underlying stock index.) Thus, a call writing strategy for the S&P 500 index should lose money on average, as it does here. For the other three markets, there is no presumption that the expected change will be either positive or negative, so our prior expectation is that call writing with reinvestment of the proceeds at the riskless rate should break even on average.

One thing that is clear is that without hedging, standard deviations are very large, and it does not make much difference whether the volatility is known or just forecasted. Indeed, since the strategy simply amounts to taking a directional bet that the underlying will not rise too far over the option's lifetime, the results are dominated by what these markets actually did during this period. In the top portion of the table, the percentage of options that ended up in the money (requiring a cash payout by the writer at maturity) for most cases was about 50% or slightly under, except for the stock index contracts. For the longer horizons in the bottom panel, percentages were lower, and results of writing at the money calls turned out a little better than could be expected. The exception is the stock market, in which the impact of the long bull market period of the 1980s and 1990s is apparent. However, the worst cases were very large losses. Some individual contracts ended up costing the writer 5 to 10 times their initial prices, and average losses in the worst years were several hundred percent.

Given the clear danger of simply writing options and holding the short position until maturity, regular writers of options tend to hedge their positions as described in Section II. Many rely on hedging delta alone, however, and not the entire range of Greek letter risk exposures. One reason for this is lack of available options that can be purchased at reasonable prices to hedge gamma risk.

[10] This is not quite a true out-of-sample test, since only at the end of the period would one know which variant produced the most accurate forecasts. These results present a best case.

Delta hedging for options based on stock indexes or exchange rates can be done by taking a position in the underlying index portfolio or currency. However, hedging options based on interest rates is normally done using futures, Eurodollar or Treasury bond futures in this case. That analysis is still in progress at the present time. A complete examination of model risk faced by a financial institution writing options and hedging its positions for all four of these markets will be presented in a subsequent paper.

Table 5 shows the results of writing at the money calls on the S&P 500 index and the Japanese yen and delta hedging them over their lifetimes. When an option is first sold, we assume its market risk is hedged by trading delta units of the underlying. Funds needed for the purchase are assumed to be borrowed at the risk free interest rate so that the option writing strategy is entirely self-financing. Subsequently, as time elapses and the market moves, each period (either every month or every day) a new delta is computed for the option using the volatility estimate based on data up to that point. The hedge is then rebalanced by adjusting the position in the hedge instrument appropriately. All cash inflows are invested and outflows are financed at the current riskless interest rate. At expiration, the hedge position is liquidated and any payoff on options finishing in the money is made. Option valuation models are based on the principle that such a trading strategy is riskless (if it could be followed continuously) and should produce a return of zero (because there is no net investment of capital--all funds required are borrowed at the riskless rate). Note that we are taking no account of transactions costs that would be incurred in pursuing a hedging strategy that can involve a large amount of trading.

The contrast with the previous table clearly shows the value of a delta neutral hedging strategy. Mean returns are relatively small and the standard deviations have been sharply reduced. Even so, there are some very bad single events, and some fairly bad full years.

These results allow us to isolate the impact of volatility forecasting errors. Inaccurate volatilities will introduce risk to the hedged call writing strategy both because options are mispriced when they are written and also because incorrect deltas will be calculated as the positions are hedged. The effect of estimation error shows up here in substantially lower standard deviations and less bad "worst" return figures when the realized volatility is used in the model in place of the forecast. However, the fact that mean returns are negative using realized volatilities indicates the presence of errors in the models themselves. This is consistent with the existence of fat tails in the probability distributions, that expose option writers to greater risk of having to make a large payout than the lognormal accounts for. It is rather unsettling to note that the standard deviation in the return for writing at the money calls and delta hedging on a daily basis using estimated volatility in the pricing model is distinctly larger than the volatility of the underlying asset itself.

III.3 Incorrect Model Implementation

Model risk can occur because of deficiencies in the model or because of inaccurate parameter values. Another real, though seemingly avoidable, source of model risk is simply errors in implementing it. Programming errors or errors in running a model (for example, entering input values incorrectly) of course should not happen. However, of course, they do happen. Implementation problems are more likely to occur for derivatives than for fundamental securities for two reasons. First, the models themselves are so much more complex that it can be very hard to find all of the programming bugs in a valuation system. Moreover, the models are called upon to do much more for derivatives than simply pricing them. A programming mistake in the calculation of gamma, for example, might lead to serious losses even if the option value is correct. Second, because the instruments are complex, pricing and hedging errors can be hard to detect with the naked eye.[11]

Although "war stories" abound, in the form of large losses at one firm or another that are attributed to incorrect pricing models, quantitative data on implementation risk is not easy to find.[12] However, in one published controlled experiment, Marshall and Siegel [1997] presented the identical asset portfolio to a number of commercial vendors of software for Value at Risk calculations. Each was asked to use the same volatility inputs, obtained from J.P. Morgan's RiskMetrics system and to report an aggregate VaR for the entire portfolio and separate figures for specific components, like swaps, and caps and floors. The variation across vendors in results were striking, even though they were all supposed to be analyzing the same position with the same methodology and parameter values. For the whole portfolio, the estimates ranged from $3.8 million to $6.1 million, and for the portion containing options, the VaR estimates varied from about $747,000 to $2,100,000.

III.4 Model Values versus Market Prices

The sources of model risk we have discussed to this point are all of a particular kind: discrepancies between the output of a particular model being used in derivatives valuation and what the "true" model with exact parameter inputs would produce. Model risk so far has come from not having the right model, not having the exact parameters for the model, or not

---

[11] This author has personal experience with a major securities clearing firm whose model for valuing foreign currency forwards allowed a trader-client to run a large unrecognized deficit in his account for more than a year. The back office staff running the model were not aware that the computer program they had obtained from an outside vendor required them to input current interest rates, and did not notice that the values for account equity it was producing were incorrect. The trader noticed, however, and quickly learned how to exploit the mistake to generate false equity in his account--which he subsequently lost in trading.

[12] In a recent issue of Risk magazine, Paul-Choudhury [1997] cites several recent losses attributed to incorrect models.

implementing the model correctly.

There is a much deeper model risk problem that we have not yet mentioned. Modern derivatives pricing models are virtually all based on the principle that the derivative can be priced quite exactly relative to its underlying asset because there is an arbitrage trade that will produce riskless excess returns if the market price for the derivative differs from the model value. For example, there is model risk in pricing options with the Black-Scholes formula because the arbitrage trade does not work the same with a fat-tailed distribution as under lognormality. A "true" model, incorporating the actual returns distribution would produce the correct price, based on a fat-tailed arbitrage strategy.

The problem is that the arbitrage strategies that these models are derived from are generally hard to execute under the best of circumstances, and actually impossible in many cases. Replicating an option payoff by arbitrage requires continuously rebalancing a hedged position, which can not be done in practice. Transactions costs from even moderately frequent rebalancing can become very large, and in many cases, the underlying can not be separately traded at all. Consider the problems of delta hedging an option on the Federal funds rate; or the call option component of a callable bond or mortgage-backed security; or a futures contract based on a stock index containing illiquid stocks or stocks with restrictions on ownership. Without arbitrage, even knowing the true model and exact input parameters does not guarantee that the market price will be close to the model value.

A pricing model gives a theoretical value for a derivative instrument, but all transactions have to be made in the market, at market prices. A major source of model risk is simply that on the date one needs to trade, the market may not obey the model, and the arbitrage trading that should theoretically force it into conformance is not strong enough to do so. This can pose great problems for valuing positions in complex derivative instruments, such as the more exotic mortgage-backed securities. Since they are not very liquid, market prices may not be available, so positions will be "marked to model" rather than marked to market.[13] But if a position must be liquidated, the market prices at which trades can be executed may be very far from the model values. An excellent example of this problem occurred when Askin Capital Management, holding a portfolio of the most complex mortgage-backed securities was forced to liquidate quickly in an unaccomodating market environment. The total loss realized on the portfolio was reported to be close to $600 million. The valuation models used in pricing these instruments may well have been correct, in terms of the arbitrage-based pricing paradigm, but the reality was that when the securities had to be sold quickly, there were no arbitrageurs prepared to buy them at their theoretical values, and the market, as always, priced them according to supply and demand.

Further evidence on this issue comes from a study of an auction of mortgage-backed derivatives

[13] See Beder [1994], for example.

by Bernardo and Cornell [1997]. A holder of a portfolio of CMOs all issued by government agencies wished to liquidate them. A public bidding process was set up that attracted broad participation from major securities firms. Although this was a public auction rather than a forced sale, and the securities involved were Federal agency issues, the range of bids received was extremely wide, reflecting a wide range of valuations across the dealer community. For the average security sold, the high bid was 63 percent above the low. Such disparity across dealer valuations indicates that whatever the market price for such derivatives is, it will be far away from many of the dealers' model values.

To the extent that many derivatives are hard to value and are traded with less liquidity among fewer market participants than fundamental securities, the problem of deviations between market prices and model values is a real risk that particularly applies to derivatives.

IV. Conclusion

Innovation in the financial markets in recent years has been spectacular, with much of it coming in the broad area encompassed by the term derivatives. Derivatives contribute to the overall strength of the financial markets by facilitating the management of financial and economic risks. A wheat futures market improves overall economic efficiency by permitting those involved in the production, storage, and distribution of wheat to hedge the price risks inherent in their activities. Hedging does not eliminate the risk--that is a product of the vagaries of weather and world grain markets and must be borne somehow by the whole economy--but it permits redistributing the burden of risk bearing away from those upon whom it naturally falls, to others who are not directly tied to the physical grain trade, but are more numerous and more able to bear risk. In the same way, the development of newer derivatives markets based on financial instruments are creating multiple channels for repackaging and redistributing much larger risk exposures of all varieties, making the overall economy better able to manage them efficiently and to withstand the effects of economic and financial shocks.

Beyond the risk shifting that new derivatives markets permit, the technology of option pricing has been extensively developed to create new types of financial instruments with payoff patterns that are especially desired by investors and other market participants. Equity investments with downside protection, loan agreements with interest rates that adjust to market conditions but can not float too high or too low, credit derivatives that allow hedging of a bond's credit risk independently of its exposure to interest rate movements, among a great many others, all contribute to the efficiency of the financial system in channeling funds from savers to borrowers. Investors can be offered securities structured with terms they find most attractive, and borrowers can receive funds under terms they find least burdensome, because of the ability of derivative instruments to transform the risk characteristics of the underlying assets that are involved.

But along with their rapid growth has come concern about the risks these new instruments entail.

Although derivatives are exposed to the same kinds of risks as more traditional securities, exactly how those risks affect them can be quite different. In this paper we have examined the major classes of financial risk and the particular ways in which they are manifested in derivative instruments.

One major type of derivatives-related risk is largely new to the financial markets: model risk. Trading in derivatives involves heavy use of complex mathematical models that are needed to understand valuation relationships and risk exposures. These models require a compromise between realism and tractability . The difference between reality and a model leads to model inaccuracy because we do not understand everything that is relevant to valuation in the real world. The need for tractability introduces model inaccuracy because it limits our ability to incorporate features of financial markets that are known to be important but are hard to model. In other words, derivatives are exposed to model risk because we don't know everything that determines true values and we can't capture everything we do know in usable models.

In this paper, I have focused on several major sources of model risk in derivatives, including the use of the lognormal probability distribution for risk assessment when actual security returns exhibit fat tails, i.e., more extreme events than the lognormal allows for, and the need to implement models with forecasted rather than true values for volatility and other critical input parameters. These sources of risk are widely known among market participants; what this paper contributes is an indication of how large their impact can be in practice, as well as a methodology that can be used to assess model risk in other cases.

In closing, then, based on the above analysis I offer these suggestions for dealing with model risk in derivatives.

1. Be aware of it. Model risk is inherent in the use of theoretical models in trading derivatives. Prudence dictates avoiding placing undue faith in model values, and being especially aware of the sources of model inaccuracy and the situations in which can be expected to have the biggest effect.

2. Estimate model risk quantitatively. Model performance can be simulated on historical data, taking care at each point that one is conducting a true out-of-sample test, using only data that would have been available to a model user at the time. Stress testing should involve examining possible inaccuracies in the valuation models themselves, both as to parameter values and model structure. It is not sufficient to conduct stress tests by assuming one's model is exactly correct and simply using it to estimate the impact of changes in market variables on the value of a given security position within the model.

3. Build formal treatment of model risk into overall risk management procedures. One strong result found in studies of model risk is that simple but robust models tend to work better when used out of sample than more ambitious but fragile ones. Ideally, valuation and risk management

models should be built that explicitly incorporate empirical, and perhaps subjective, estimates of parameter uncertainty and model drift. Along these lines, there is a role for both model-based and non-model based techniques in overall risk management. Finally, it makes sense to recognize cases in which one's model is weak and limit the firm's exposure in markets for which sturdy models are not available.

Appendix

The original Black and Scholes [1993] model for option pricing applies to a European call option on a non-dividend paying stock. This model can easily be modified to apply to an underlying asset that makes a regular cash payout at a continuous rate q during the life of the option. The version used for stock index options sets q equal to the rate of dividend payout on the underlying index portfolio. For options on foreign currencies, q is the foreign riskless interest rate, which leads to the Garman-Kohlhagen [1983] formula. For options on interest rates (and futures contracts), q is set equal to the riskless interest rate r, giving the formula presented in Black [1976].

The general valuation equation for a call option is given by equation (A.1).

$$C = Se^{-qT} N[d] - Xe^{-rT} N[d - \sigma\sqrt{T}] \qquad \text{(A.1)}$$

where

$$d = \frac{\ln\frac{S}{X} + (r - q + \frac{\sigma^2}{2})T}{\sigma\sqrt{T}}$$

and C is the call value, S is the current price or level of the underlying, X is the strike price, T is the time to option maturity, r is the (domestic) riskless interest rate, q is the appropriate value for the option in question as defined above, and σ is the volatility of the underlying.

The formula for a put option, P, is given by

$$P = Xe^{-rT} N[-d + \sigma\sqrt{T}] - Se^{-qT} N[-d] \qquad \text{(A.2)}$$

References

Barron's. 77, No. 44, Nov. 3, 1997, p. MW 17.

Basle Committee on Banking Supervision. Risk Management Guidelines for Derivatives. Basle: Bank for International Settlements, 1994.

Basle Committee on Banking Supervision. Amendment to the Capital Accord to Incorporate Market Risks. Basle: Bank for International Settlements, 1996.

Beder, Tanya Styblo. "The Realities of Marking to Model." Journal of Bank Accounting and Finance , 1994, pp. 4-12.

Bernardo, Antonio and Bradford Cornell. "The Valuation of Complex Derivatives by Major Investment Firms: Empirical Evidence." Journal of Finance 52, June 1997, pp. 785-798.

Black, Fischer. "The Pricing of Commodity Contracts." Journal of Financial Economics 3, 1976, pp. 167-179.

Black, Fischer and Myron Scholes. "The Pricing of Options and Corporate Liabilities." Journal of Political Economy 81, 1973, pp.637-659.

Canina, Linda and Stephen Figlewski. "The Informational Content of Implied Volatility." Review of Financial Studies 6, 1993, pp. 659-681.

Figlewski, Stephen. "Financial Derivatives: Actions Needed to Protect the Financial System: The Government Accounting Office Report on Derivatives Regulation: A Synopsis and Commentary." Journal of Derivatives 2, Fall 1994, pp. 80 - 93.

Figlewski, Stephen. "Forecasting Volatility." Financial Markets, Institutions & Instruments 6, 1997.

Garman, Mark B. and Steven W. Kohlhagen. "Foreign Currency Option Values." Journal of International Money and Finance 2, 1983, pp. 231-237.

Group of Thirty Global Derivatives Study Group. Derivatives: Practices and Principles. Washington: Group of Thirty, 1993.

Marshall, Chris and Michael Siegel. "Value at Risk: Implementing a Risk Measurement Standard." Journal of Derivatives 4, Spring 1997, pp. 91-111.

Risk Standards Working Group. Risk Standards for Institutional Investment Managers and

Institutional Investors. 1996.

Sharpe, William F. Investments, 3rd ed. Englewood Cliffs, NJ: Prentice-Hall, 1985.

Smith, Donald J. "Aggressive Corporate Finance: A Close Look at the Procter & Gamble-Bankers Trust Leveraged Swap." Journal of Derivatives 4, Summer 1997, pp. 67-79.

Paul-Chaudhury, Sumit. "This Year's Model." Risk 10, May 1997, pp. 18-23.

U.S. Government Accounting Office. Financial Derivatives: Actions Needed to Protect the Financial System. Washington: U.S. Printing Office, 1994.

Figure 1: Market Risk Exposure in Trading Options

Writing Japanese Yen Calls with $100 Initial Value

S_0 = 90, X = 90, T = 3 months, Vol'y = 0.12, r = 6%, ryen = 2%

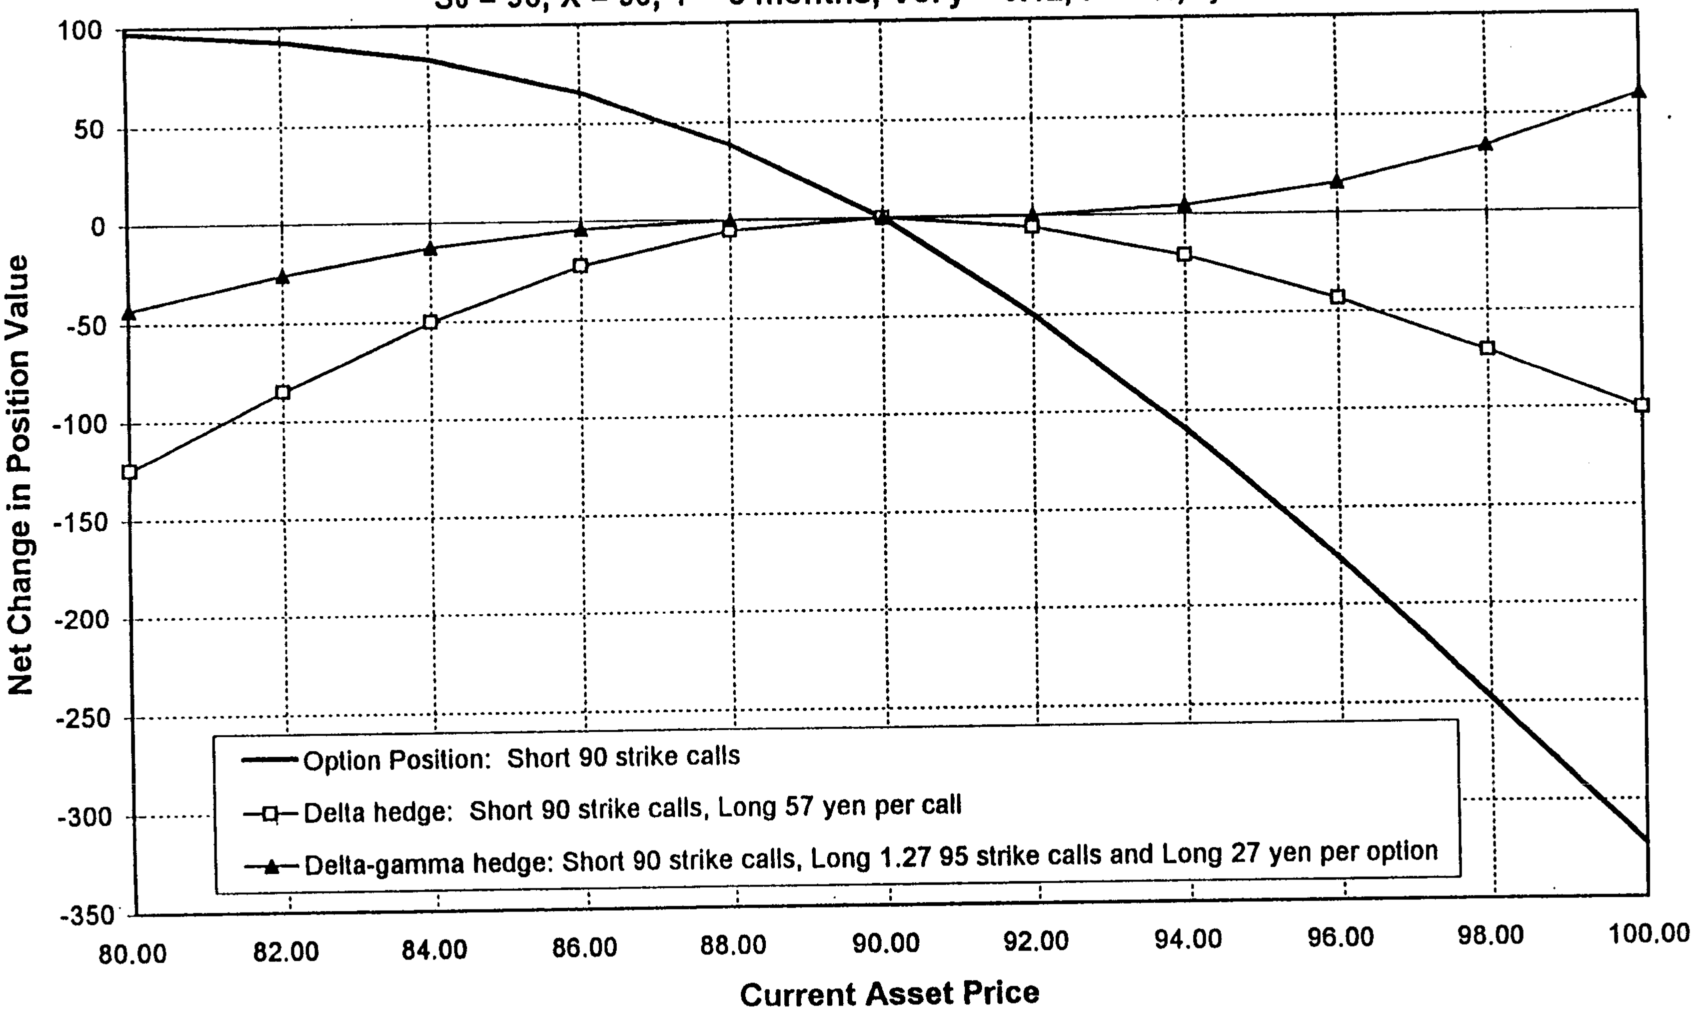

Figure 2: Market Risk Exposure in Trading Options

Delta-Gamma Hedge of Written Japanese Yen Calls as Volatility Changes

$S_0 = 90$, $X = 90$, $T = 3$ months, Original Vol'y = 0.12, $r = 6\%$, $r_{yen} = 2\%$

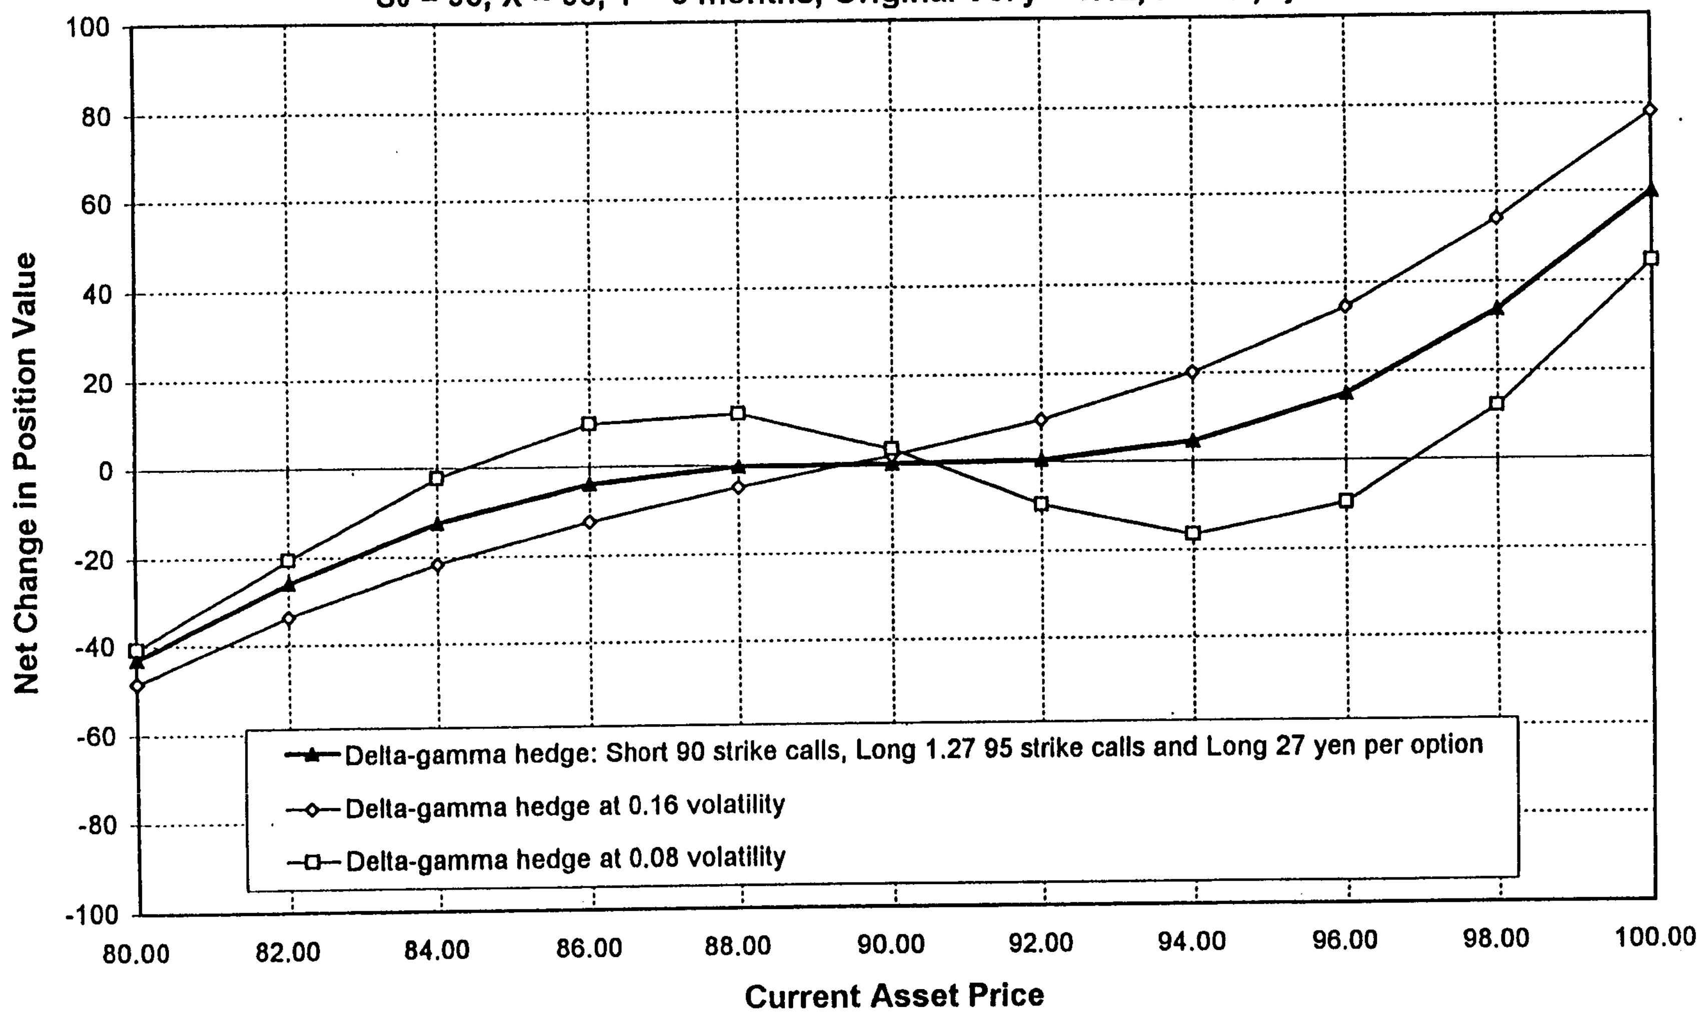

Figure 3

Time Pattern of Credit Risk Exposure for Interest Rate Derivative Contracts

Face value = $1 million; Interest rate = 6-month LIBOR; Semiannual repricing for swap and cap

Maturity = 3 years; r_0 = 6.00; Strike rate = 6.00; Volatility = 20%

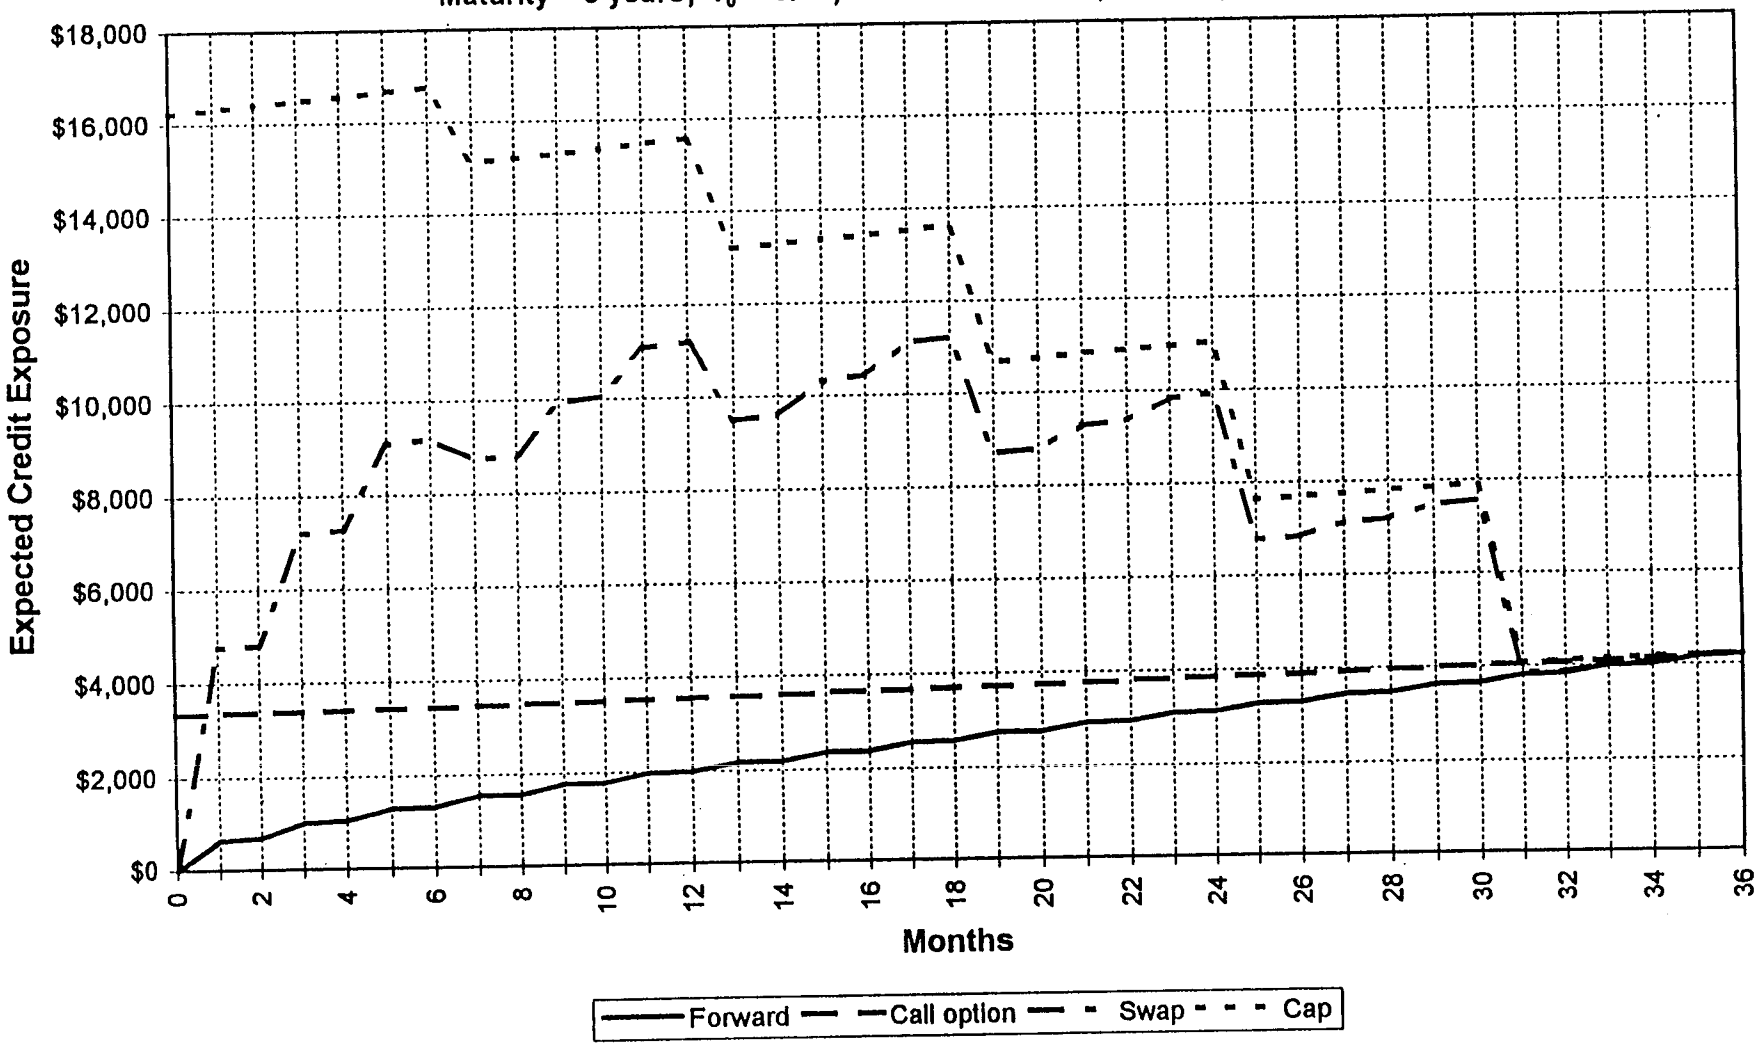

Figure 4: Fat-Tailed versus Normal Distribution

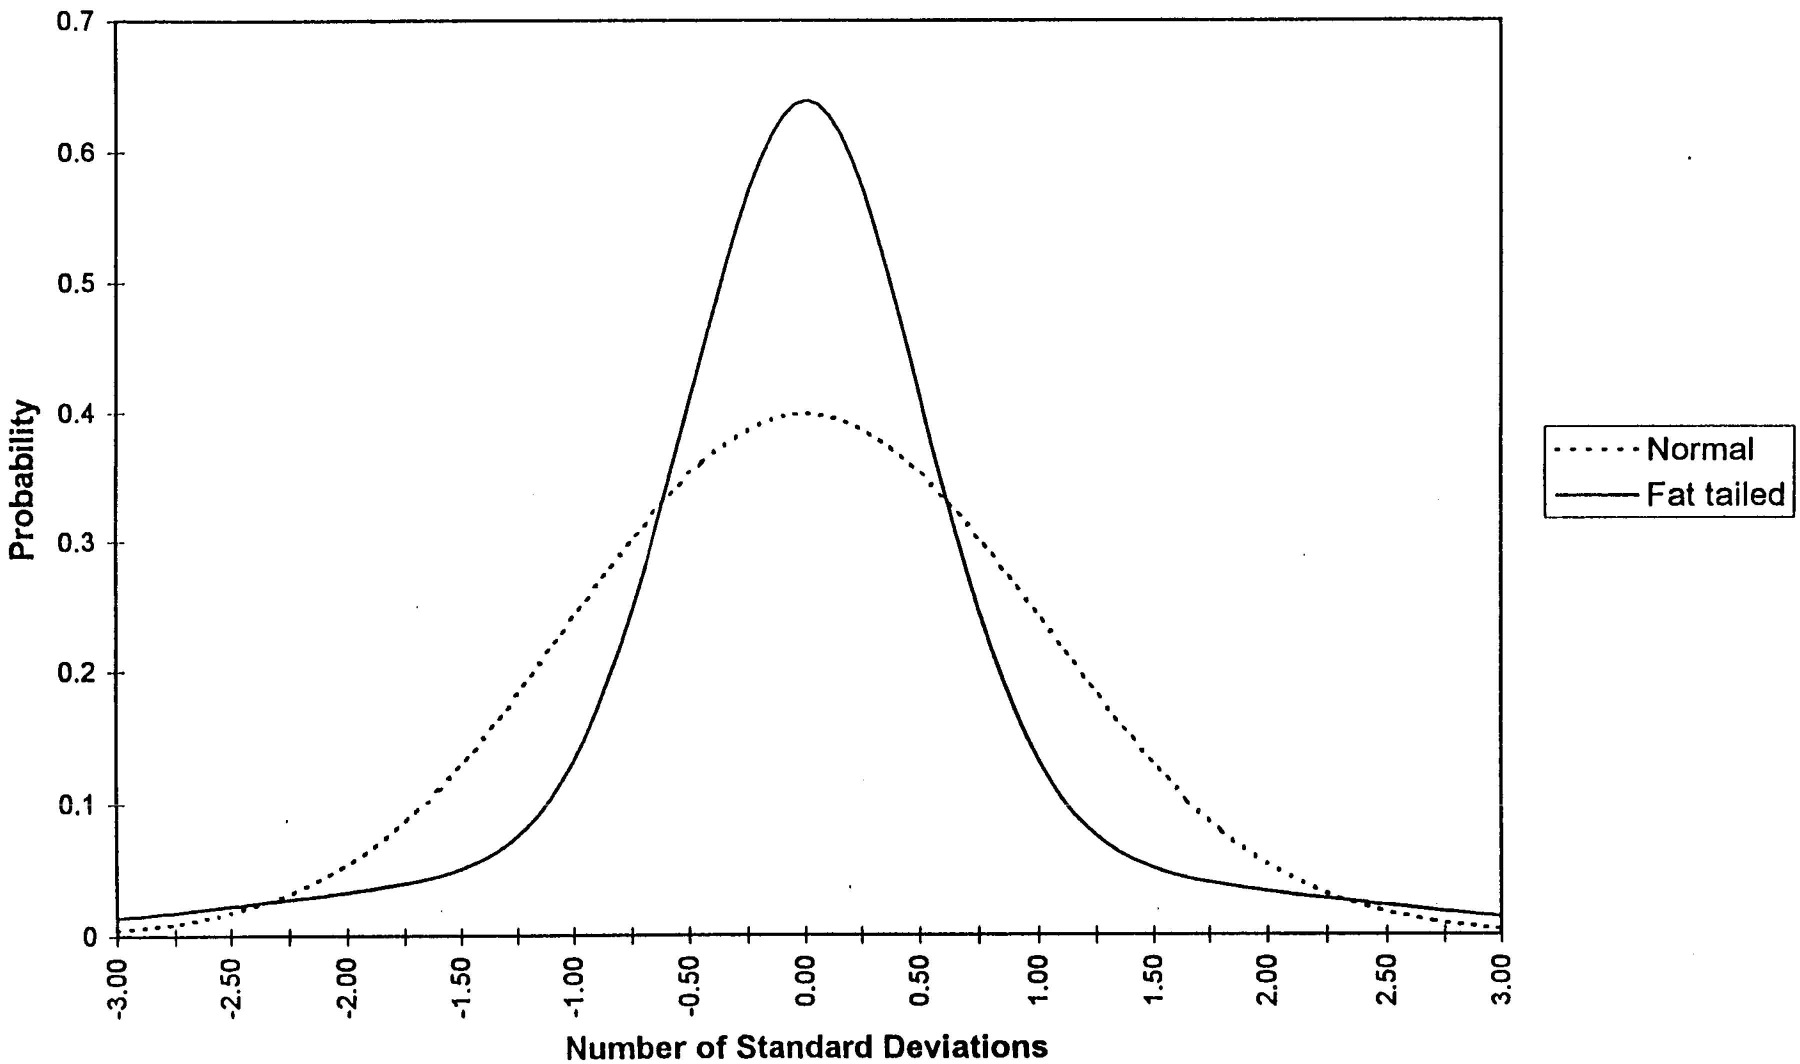

Table 1

Model Estimates of Risk Exposure for Short Yen Call Position and Two Hedges

A. Model Values

Parameters: S = 90, X = 90, T = 3 months, r = 6.0%, r_{YEN} = 2.0%, σ = 0.12

Call Value	2.57
Delta	0.57
Gamma	0.073
Lambda	20.1
Theta	-0.017
Vega	0.174
Rho	0.001

Note: All values are in U.S. cents per 100 yen. Delta is the change in option value per unit change in S; Gamma is the change in delta per unit change in S; Lambda, the leverage ratio or elasticity, is the percent change in option value per percent change in the underlying; Theta is the change in option value over 1 day if S is unchanged; Vega is the change in option value for .01 increase in volatility; Rho is the change in option value for 1 basis point increase in r (holding S fixed).

B. Position Risk Exposure

All positions involve writing \$100 worth of 3 month 90 strike call options. The "Delta Hedge" position hedges market risk by holding 0.57 units of the underlying (i.e., 57 yen) per call option written. The "Delta-Gamma Hedge" position hedges gamma risk by purchasing 1.27 3-month 95 strike calls at the initial model price of 0.718, and overall delta risk by purchasing 27 yen, per call option written. The table shows the model estimate for the change in position value over 1 day, for various ending spot exchange rates and volatility levels.

Position	Volatility	S = 80	S = 85	S = 90	S = 95	S = 100
Unhedged	0.12	97.4	76.2	0.6	-141.7	-322.2
Delta Hedge						
	0.12	-124.4	-34.6	0.8	-30.6	-100.2
	0.08	-122.0	-18.2	27.6	-18.5	-98.1
	0.16	-131.9	-55.5	-26.3	-48.9	-107.4
Delta-Gamma Hedge	0.12	-42.9	-7.2	0.1	8.6	59.6
	0.08	-40.7	4.9	3.6	-15.2	44.3
	0.16	-48.2	-16.4	1.9	26.6	77.8

Table 2: Fat Tailed Returns Distributions and Errors in Assessing Risk Exposure for Several Markets

	Lower 5% Tail				Upper 5% Tail			
	Tail cutoff: -1.64 std dev; Predicted mean in 5% tail: -2.06 std dev				Tail cutoff: 1.64 std dev; Predicted mean in 5% tail: 2.06 std dev			
	% of returns in predicted 5% tail	Mean actual return in predicted 5% tail	Cutoff for actual 5% tail	Mean return in actual 5% tail	% of returns in predicted 5% tail	Mean actual return in predicted 5% tail	Cutoff for actual 5% tail	Mean return in actual 5% tail
3 Month LIBOR	2.80%	-3.42	-1.15	-2.50	2.20%	3.82	1.15	2.44
10 Year T-bond Yield	5.80%	-2.21	-1.79	-2.29	3.60%	2.87	1.51	2.51
S&P 500 Index	4.00%	-2.75	-1.50	-2.51	7.80%	2.09	1.80	2.29
Yen Exchange Rate	3.60%	-2.65	-1.45	-2.34	4.80%	2.67	1.62	2.63

	Lower 1% Tail				Upper 1% Tail			
	Tail cutoff: -2.33 std dev; Predicted mean in 1% tail: -2.66 std dev				Tail cutoff: 2.33 std dev; Predicted mean in 1% tail: 2.67 std dev			
	% of returns in predicted 1% tail	Mean actual return in predicted 1% tail	Cutoff for actual 1% tail	Mean return in actual 1% tail	% of returns in predicted 1% tail	Mean actual return in predicted 1% tail	Cutoff for actual 1% tail	Mean return in actual 1% tail
3 Month LIBOR	1.60%	-4.52	-3.25	-5.70	1.00%	6.29	3.00	6.29
10 Year T-bond Yield	1.80%	-2.90	-2.67	-3.25	2.00%	3.63	3.62	4.43
S&P 500 Index	2.60%	-3.15	-3.28	-3.99	1.40%	2.84	2.75	2.96
Yen Exchange Rate	2.00%	-3.15	-3.35	-3.78	2.60%	3.35	3.70	4.13

NOTE: For 500 days, from Jan 12, 1995 to Dec. 31, 1996 a volatility estimate for each day was calculated from the previous 250 days' data. The realized price or yield change for that day was compared to the predicted tail cutoff values from a lognormal distribution with volatility equal to the estimated value. The table presents results in terms of standard deviations. To convert standard deviations approximately to yield or percentage price changes, multiply by the average 1 day estimated volatilities for the period, as follows: Average 1 day standard deviations: 3 Month LIBOR = 5.14 b.p.; 10 Year T-bond Yield = 6.03 b.p.; S&P 500 Index = 0.599%; Yen Exchange Rate = 0.74%

Table 3a

Forecast Accuracy of Volatilities Estimated from Historical Data
Daily Data

The table shows root mean squared forecast error for annualized volatility calculated from daily data around a mean of zero, for different forecast horizons and historical sample lengths. Also reported is the performance of exponential weighting across all past data, in which the optimal weight is chosen to minimize the forecast error for each forecast horizon. One "month" is defined as 21 trading days. Shading indicates minimum RMSE sample size.

London Inter-Bank 3 Month Rate
November 1, 1979 - January 12, 1996

Months in Sample	Forecast Horizon (Months) 1	3	12
3	10.2%	9.3%	8.4%
12	10.2%	8.8%	8.1%
60	13.0%	11.5%	10.3%
Exp. wgt'ed	9.8%	8.7%	8.0%
Avg. wgt.	0.973	0.984	0.987
Average Realized	20.8%	21.5%	21.5%

10 Year Treasury Bond Yield
January 26, 1976 - January 5, 1996

Months in Sample	Forecast Horizon (Months) 1	3	12
3	4.6%	4.4%	4.4%
12	4.8%	4.2%	4.0%
60	5.7%	5.1%	4.6%
Exp. wgt'ed	4.7%	4.6%	4.7%
Avg. wgt.	0.976	0.986	0.992
Average Realized	12.7%	13.0%	13.5%

S&P 500 Stock Index
December 30, 1975 - January 3, 1996

Months in Sample	Forecast Horizon (Months) 1	3	12
3	7.2%	7.2%	7.1%
12	7.5%	7.1%	6.8%
60	7.8%	7.2%	6.5%
Exp. wgt'ed	7.0%	6.7%	6.5%
Avg. wgt.	0.979	0.986	0.990
Average Realized	12.9%	13.2%	13.7%

Yen Exchange Rate
January 30, 1976 - January 4, 1996

Months in Sample	Forecast Horizon (Months) 1	3	12
3	3.7%	3.2%	3.1%
12	3.6%	3.0%	2.6%
60	3.8%	3.2%	2.5%
Exp. wgt'ed	3.7%	3.1%	2.5%
Avg. wgt.	0.986	0.991	0.995
Average Realized	9.5%	9.8%	10.1%

Table 3b

Forecast Accuracy of Volatilities Estimated from Historical Data
Monthly Data

The table shows root mean squared forecast error for annualized volatility calculated from monthly data around a mean of zero, for different forcast horizons and historical sample lengths. Also reported is the performance of exponential weighting across all past data, in which the optimal weight is chosen to minimize the forecast error for each forecast horizon. Shading indicates minimum RMSE sample size.

London Inter-Bank 3 Month Rate
Jan 1976 - Dec 1991

Months in Sample	Forecast Horizon 24	60
24	13.3%	11.9%
60	12.6%	10.5%
All available	15.2%	13.5%
Exp. wgt'ed	13.7%	12.8%
Avg. wgt.	0.898	0.942
Average Realized	25.4%	25.3%

10 Year Treasury Bond Yield
Jan 1976 - Dec 1991

Months in Sample	Forecast Horizon 24	60
24	5.4%	5.0%
60	4.7%	4.2%
All available	4.4%	3.4%
Exp. wgt'ed	5.3%	3.7%
Avg. wgt.	0.971	0.991
Average Realized	15.2%	15.7%

S&P 500 Stock Index
Jan 1976 - Dec 1991

Months in Sample	Forecast Horizon 24	60
24	5.9%	5.4%
60	4.9%	4.5%
All available	4.0%	3.2%
Exp. wgt'ed	4.2%	4.1%
Avg. wgt.	0.976	0.987
Average Realized	15.4%	15.2%

Yen Exchange Rate
Jan 1976 - Dec 1991

Months in Sample	Forecast Horizon 24	60
24	4.3%	4.1%
60	3.6%	2.8%
All available	3.3%	2.7%
Exp. wgt'ed	3.6%	2.9%
Avg. wgt.	0.957	0.982
Average Realized	12.1%	12.1%

Table 4

Return and Risk in Writing Options with Estimated Volatilities: At the Money Calls

The table reports the performance of a strategy of writing unhedged at the money call options each period for several maturities. The period is daily in the upper panel and monthly in the lower. Option prices are computed by using as the volatility input either the minimum RMSE volatility estimate using historical data (left panel) or the realized volatility over the option's lifetime. The strategy always writes enough options to equal $100 of option premium. The results can therefore be interpreted either as dollars or as percentages of the initial option price.

		Minimum RMSE Volatility Forecast using Historical Data							Realized Volatility						
Series	Option Life	Mean Return	Standard Deviation	% ITM	Worst Case	Worst Return	Worst Year	Worst Mean	Mean Return	Standard Deviation	% ITM	Worst Case	Worst Return	Worst Year	Worst Mean
Daily Data															
S&P 500 Stock Index	1 mon	-19.36	158.08	0.60	8/12/82	-1035.20	1975	-525.52	-26.89	155.89	0.60	1/15/91	-627.15	1975	-528.51
	3 mon	-29.35	158.59	0.67	8/11/82	-1077.77	1975	-334.20	-35.57	155.19	0.67	8/11/82	-606.38	1975	-387.99
	1 yr	-56.22	164.20	0.76	7/28/82	-688.87	1982	-344.38	-78.25	173.38	0.76	12/8/94	-706.15	1995	-340.70
3-mo. LIBOR	1 mon	11.23	144.73	0.45	2/1/94	-1033.78	1994	-160.26	13.79	129.35	0.44	4/8/94	-633.96	1994	-139.82
	3 mon	14.63	144.48	0.48	9/15/80	-730.76	1994	-249.73	11.35	142.82	0.45	9/3/80	-620.07	1994	-249.23
	1 yr	16.28	177.62	0.44	12/17/93	-916.42	1993	-328.53	10.74	186.30	0.40	1/11/94	-931.51	1993	-342.60
Ten-Year Treasury Yield	1 mon	4.02	154.50	0.49	9/21/79	-1191.34	1987	-93.40	9.00	128.38	0.49	9/21/79	-525.12	1994	-54.38
	3 mon	16.55	129.54	0.52	2/3/94	-699.08	1979	-114.20	17.81	118.91	0.52	7/17/87	-503.14	1994	-70.84
	1 yr	32.91	109.75	0.50	10/20/93	-551.04	1993	-209.03	34.64	103.54	0.50	10/20/93	-459.06	1993	-169.41
Yen Exchange Rate	1 mon	15.05	136.22	0.47	9/30/76	-953.09	1995	-40.27	9.76	136.28	0.47	2/7/92	-647.68	1979	-54.01
	3 mon	22.55	128.24	0.42	6/20/95	-687.35	1995	-108.41	18.90	131.14	0.42	6/20/95	-572.48	1995	-84.80
	1 yr	61.00	93.05	0.35	10/30/78	-404.50	1989	-46.94	62.35	88.22	0.31	10/30/78	-349.58	1981	-40.19
Monthly Data															
S&P 500 Stock Index	2 yr	-59.29	162.65	88%	Sep-85	-602.51	1985	-369.70	-58.39	160.53	88%	Sep-85	-528.10	1984	-320.26
	5 yr	-173.43	215.64	100%	Jul-82	-883.74	1982	-568.82	-163.13	206.69	100%	Jul-82	-874.96	1982	-562.95
3-mo. LIBOR	2 yr	32.69	134.30	47%	Mar-78	-470.70	1977	-239.82	21.57	158.80	47%	Dec-76	-543.08	1977	-322.37
	5 yr	119.65	105.32	23%	Aug-76	-293.05	1976	-200.60	119.47	106.92	23%	Aug-76	-309.92	1976	-214.30
Ten-Year Treasury Yield	2 yr	53.22	82.13	55%	Feb-78	-176.35	1979	-69.20	54.91	79.26	55%	Feb-78	-158.33	1979	-49.48
	5 yr	119.06	73.34	44%	Sep-76	-96.13	1976	-49.51	119.96	72.07	44%	Sep-76	-91.68	1976	-45.64
Yen Exchange Rate	2 yr	83.08	64.72	41%	Apr-88	-192.25	1988	-43.95	85.89	59.82	41%	Apr-88	-172.47	1988	-34.72
	5 yr	152.61	33.62	29%	Oct-78	6.24	1978	99.89	153.79	31.06	29%	Oct-78	21.36	1978	107.26

Table 5

Return and Risk in Writing and Delta Hedging Options with Estimated Volatilities: At the Money Calls

The table reports the performance of a strategy of writing and hedging at the money call options each period for several maturities. The period is daily in the upper panel and monthly in the lower. Option prices are computed by using as the volatility input either the minimum RMSE volatility estimate using historical data (left panel) or the realized volatility over the option's lifetime. The hedge is rebalanced to be delta neutral every period, using the current volatility estimate. The strategy always writes enough options to equal $100 of option premium. The results can therefore be interpreted either as dollars or as percentages.

		Minimum RMSE Volatility Forecast using Historical Data							Realized Volatility						
Series	Option Life	Mean Return	Standard Deviation	% ITM	Worst Case	Worst Return	Worst Year	Worst Mean	Mean Return	Standard Deviation	% ITM	Worst Case	Worst Return	Worst Year	Worst Mean
Daily Data															
S&P 500	1 mon	0.18	33.32	0.60	10/16/87	-536.22	1982	-24.51	-3.19	12.39	0.60	8/12/86	-80.93	1986	-5.07
Stock Index	3 mon	1.99	24.62	0.67	10/16/87	-329.05	1986	-23.73	-2.53	5.34	0.67	6/12/86	-45.34	1980	-5.25
	1 yr	3.26	42.40	0.76	1/9/87	-290.53	1987	-99.22	-4.64	3.82	0.76	10/27/86	-54.66	1980	-10.75
Yen	1 mon	0.75	38.04	0.47	3/27/78	-389.84	1978	-19.57	-2.08	14.65	0.47	5/7/76	-82.38	1981	-7.48
Exchange Rate	3 mon	-0.83	24.54	0.42	3/24/78	-137.27	1978	-20.41	-2.36	5.58	0.42	11/15/93	-59.74	1981	-5.55
	1 yr	1.54	19.32	0.31	9/5/85	-57.30	1979	-21.12	-2.48	3.51	0.35	5/25/79	-16.38	1981	-9.31
Monthly Data															
S&P 500	2 yr	-7.50	18.58	0.88	Jun-86	-57.17	1986	-43.87	-6.41	10.80	0.90	Aug-85	-23.48	1982	-14.13
Stock Index	5 yr	-14.06	6.56	1.00	Sep-78	-22.88	1979	-22.10	-11.85	9.96	0.99	Oct-78	-21.89	1978	-21.07
Yen	2 yr	-11.79	16.18	0.41	May-78	-48.00	1978	-31.48	-2.44	9.58	0.38	May-76	-23.97	1983	-9.40
Exchange Rate	5 yr	-14.99	12.19	0.29	Dec-77	-55.92	1977	-36.98	-1.66	8.04	0.19	Aug-78	-9.58	1978	-6.78

Version of Dec. 7, 1997

Comments Welcome

Derivatives Risks, Old and New

by

Stephen Figlewski

Professor of Finance

New York University Stern School of Business
44 West 4th St., Suite 9-160
New York, NY 10012-1126

Phone: 212-998-0712
FAX: 212-995-4220

E-mail: sfiglews@stern.nyu.edu

I would like to thank T. Clifton Green and Sari Carp for invaluable assistance in preparing this paper. I am grateful for valuable discussions and comments on earlier drafts provided by Kenneth Froewiss, Robert Litan, Leslie Rahl, Alexander Reisz, Anthony Santomero, Charles Smithson, René Stulz and seminar participants at NYU and at the Brookings-Wharton Conference.

Abstract

There has been much discussion of risks tied to trading in derivatives, with some well-informed objective observers arguing that derivatives risks are not significantly greater or different from those associated with traditional financial instruments. Financial risks are often broken down into market risk, credit risk, operational risk and legal risk. We review the standard classification and observe that while derivatives are exposed to these types of risk, they are manifested quite differently in derivatives than in traditional securities. We then consider a "new" type of risk that is particularly important for derivatives: model risk. Derivatives trading depends heavily on the use of theoretical valuation models, but these are susceptible to error from incorrect assumptions about the underlying asset price process, estimation error on volatility and other inputs that must be forecasted, errors in implementing the theoretical models, and differences between market prices and theoretical values. Empirical evidence drawn from several important asset markets shows that model error can be quite large and can be expected to lead to significant risk in derivatives pricing and risk management.