Proceedings of the 2002 Winter Simulation Conference
E. Yücesan, C.-H. Chen, J. L. Snowdon, and J. M. Charnes, eds.

SECURITY PRICE DYNAMICS AND SIMULATION IN FINANCIAL ENGINEERING

Stewart Mayhew

Terry College of Business
University of Georgia
Athens, GA 30606–6253, U.S.A.

ABSTRACT

Applications in financial engineering have relied heavily on Brownian Motion as a workhorse model for pricing derivative securities and implementing risk management programs. When more than one state variable is required, the standard approach is to use a multivariate Brownian Motion with constant correlations. This article briefly summarizes several important reasons why this approach is not adequate (and in some cases, can lead to disaster). Examples include fat tails, volatility clustering, large discrete jumps, parameter instability, and asymmetric correlations. Including such features makes analytic modelling less tractable, and potentially makes simulation a more attractive alternative.

1 INTRODUCTION

One of the main questions facing the financial engineering profession is how best to model the dynamics of prices in financial markets, such as individual stock prices, stock indices, commodities, interest rates, and exchange rates. Armed with a model of price dynamics, the financial engineer can:

- calculate theoretical prices for options, swaps, and other derivative securities,
- measure the amount of risk associated with holding these positions,
- aggregate risk measures across a large number of positions, and
- identify offsetting positions in other securities to control aggregate risk exposure.

The textbook approach is to specify a diffusion process for each underlying asset price—that is, a stochastic integral or stochastic differential equation where the uncertainty is driven by Brownian Motion. For stocks and stock indices, the typical assumption is Geometric Brownian Motion, with a constant volatility. This implies a normal distribution for the continuously-compounded stock return, or a lognormal distribution for the stock price. Correlations between asset prices are generally assumed to be constant over time. For interest rates, the process usually includes a mean-reverting drift term, and sometimes extra state variables are added to give the models additional flexibility. But still, the underlying state variables are modelled as diffusion processes.

The widespread adoption of Brownian Motion as a framework for describing changes in financial asset prices is most likely due to its analytic tractability. Under Brownian Motion, prices of options and other simple derivative securities can often be expressed as functions of the Normal CDF, and these prices can be computed very rapidly. In a risk management context, the assumption of joint normality in asset returns makes it fairly easy to calculate the tail probabilities for portfolios of assets. Moreover, for problems too complicated to solve analytically, the assumption of constant volatility allows for a very simple implementation of a discrete-time lattice model, or a binomial tree. In addition, this framework lends itself naturally to simulation, as it is a simple matter to generate pseudo-random deviates from a multivariate normal distribution with a known covariance structure.

Unfortunately, it appears that asset prices in the real world are not driven by diffusion processes, nor are variances and covariances constant over time. There is now a vast literature on the time-series modelling of asset prices. This research has revealed many empirical failures of the constant-parameter multivariate diffusion model. In some cases, using a multivariate diffusion model can lead to large errors in estimated prices and hedge parameters.

In risk management, success or failure depends critically on the accuracy of the model specified for the joint dynamics of the underlying asset prices. To illustrate, consider the notorious case of the hedge fund Long Term Capital Management (LTCM). It appears that immediately prior to the Russian debt crisis in August 1998, LTCM had calculated their daily "Value-At-Risk" (VAR) to be $35

Million. That is, a one-day loss in excess of $35 Million would constitute a "rare event." Lowenstein (2000) reports that their VAR numbers implied that it would take a ten standard deviation rare event for LTCM to lose all their capital in one year. They were wiped out in five weeks. LTCM lost $553 million (more than 15 times VAR) on August 21, 1998, and $277 million more on August 27. The losses continued throughout September, including another loss exceeding $500 million on September 21, and many other daily losses in excess of $100 million. According to LTCM's risk management model, losses of such magnitude were virtually impossible.

In retrospect, it appears that LTCM had failed to recognize that the correlations between their various risk exposures would be much, much higher during an international crisis than during normal times. LTCM's capital was spread out across positions in various markets and strategies. Their analysis of historical time-series data led them to believe that the correlations between their individual positions were relatively low, so that the overall portfolio was thought to be well-diversified and low-risk. A rare event in one market might generate large losses to a particular position, but as each position represented only a small portion of the whole, it was thought, the portfolio would not suffer unduly large losses. When the crisis hit, LTCM suffered large losses simultaneously in nearly all its major positions.

Those who aspire to develop effective and reliable risk management tools should seek a deeper understanding of the dynamics of asset price returns. This essay briefly reviews some of the more important empirical regularities that have been documented in the literature, including fat tails (Section 2), volatility clustering (Section 3), large jumps (Section 4), non-Gaussian copulas (Section 5), and parameter instability (Section 6). A few of these features have been incorporated into theoretical option pricing models, and in some special cases closed-form solutions are available. But in general, incorporating these features tends to make the models analytically intractable. These features can also make lattice-based recursive models considerably more complicated. However, most of these features can be simulated without undue difficulty. Historically, many in the industry have been reluctant to embrace simulation due to concerns about computation speed. These concerns are growing less relevant over time, as a result of technological improvements in processor speed and parallel processing, along with variance reduction techniques and other improvements in simulation technology. It is my belief that over time, simulation will become more firmly established as the premier technique for measuring risk in financial markets.

2 FAT TAILS

It has long been known that the distribution of security returns is leptokurtic—that is, tends to exhibit "fat tails" relative to the normal distribution. Work by Mandelbrot (1963), Fama (1965) and others inspired a flurry of research activity in the 1960s and 1970s in search of alternative distributions.

Fat tails are especially pronounced in the distribution of stock returns over short horizons. To illustrate, let us consider the unconditional distribution of daily stock index returns in the United States from 2/17/1885 through 8/15/2002, a sample of 32,572 index returns. These data are available from Professor Schwert (1990) at The University of Rochester and from the Center for Research in Security Prices (CRSP). Suppose that one were to model these returns as being independent and identically distributed (*iid*) draws from a normal distribution. The sample standard deviation of percentage daily stock returns over this time period is 0.0102, and the mean is .0003. Based on a normal distribution, we would expect to see a return more than five standard deviations below the mean approximately once in 13,822 years (based on 252 trading days/year). In fact, this has occurred 66 times since 1885, most recently on April 14, 2000. On the positive side, there have been 56 five-standard deviation events. At seven standard deviations, there have been 18 negative and 22 positive events. The crash of October 19, 1987 was a 20-standard deviation event. See Table 1 for a list of the ten largest one-day percentage increases and decreases.

Table 1: Extreme Negative and Positive Returns, 1885-2002

Date	Return	Date	Return
19871019	-0.2047	19330315	0.1665
19291028	-0.1233	19291030	0.1254
19291029	-0.1015	19311006	0.1238
19291106	-0.0990	19320921	0.1183
19371018	-0.0927	19390905	0.0964
19330720	-0.0888	19330420	0.0954
19330721	-0.0870	19871021	0.0910
18951220	-0.0851	19291114	0.0897
19871026	-0.0828	19320803	0.0888
19321005	-0.0819	19311008	0.0861

Clearly, the assumption of normally distributed returns is not appropriate in this market. As a result of this excess kurtosis, option pricing models based on the lognormal assumption, such as the Black-Scholes model, will severely underprice short-term out-of-the-money options. In addition, Value-at-Risk estimates (estimates of tail probabilities) based on the normal distribution will be way too low, leading risk managers to dramatically underestimate the risk of large losses.

One alternative that gives fat tails is the bivariate diffusion stochastic volatility family of models, such as Heston's (1993) model. In these models, the instantaneous diffusion coefficient in the stock process itself follows a diffusion process. But these models also tend to underprice short-term out-of-the-money options, because the distribution gets close to normal as the time horizon shrinks. In order to model the extreme fat tails over short horizons, it appears that some type of jump model is necessary (see Bakshi, Cao and Chen 1997).

3 VOLATILITY CLUSTERING

If fat tails were the only problem, an obvious remedy would be to model returns as being drawn from a fat-tailed distribution. Over the years, various authors have suggested using the t distribution, a mixture of normals, or a member of the Stable Paretian family, among others. This approach is not fully adequate, however, as there is clear evidence that returns are not *iid*.

A clear example is the phenomenon of "volatility clustering" or "volatility persistence"—the tendency of extreme returns to cluster together in time. If returns were fat-tailed but *iid*, extreme events would tend to be evenly distributed over time—we would be no more likely to observe an extreme event the day after a stock market crash as on any other day. But in fact, we tend to observe extended episodes of high volatility.

To illustrate, imagine estimating a risk-management model on October 22, 1929. The stock market had not moved up or down five percent in a single day in nearly twelve years. One might have easily classified a five percent move in one day to be a rare event. As shown in Table 2, it occurred twelve times in the following four weeks. Prior to October 16, 1987, the market had not declined five percent in one day in more than twenty-five years (May 28, 1962), or increased five percent in more than seventeen years (May 27, 1970). There are many other examples. A few are shown in Table 2.

Table 2: Examples of Volatility Clustering

Date	Return
19291023	-0.0590
19291028	-0.1232
19291029	-0.1015
19291030	0.1254
19291031	0.0505
19291104	-0.0523
19291106	-0.0990
19291111	-0.0620
19291112	-0.0566
19291113	-0.0569
19291114	0.0897
19291115	0.0554

Date	Return
19330719	-0.04914
19330720	-0.0888
19330721	-0.0870
19330724	0.0814
19871016	-0.0516
19871019	-0.2047
19871020	0.0533
19871021	0.0910
19871026	-0.0828
19971027	-0.0687
19971028	0.0512

Volatility clustering has important implications for risk management. Risk management is most important during periods of crisis. A good measure of volatility persistence can help the risk manager assess the likelihood of an "aftershock" following a major market move, and thus can improve risk measures exactly when an accurate measure is needed most.

Volatility clustering also has direct implications for the calculation of option hedge ratios. When a stock price declines, this will generally cause the price of a call option to decline. If volatility is highly persistent, however, a sudden decline in the stock price will cause the market's forecast of subsequent volatility to increase, and this has an offsetting effect on the option price. As a result, the call option price will fall less than it would have in a world without volatility persistence.

Various time-series techniques have been proposed to account for volatility clustering. One popular approach is to use some variant of the discrete-time Generalized Autoregressive Conditional Heteroskedasticity (GARCH) model, pioneered by Engle (1982) and Bollerslev (1986). In these models, the conditional variance of today's stock return is a function of yesterday's conditional variance and of yesterday's realized return. Extreme returns increase the variance over the following days, making additional extreme returns more likely.

There are other ways to allow for short episodes of high volatility. For example, one might use a continuous-time stochastic volatility model, or a regime-switching model, in which the variance of returns randomly switches between multiple levels (see Hamilton 1994).

Further research on the conditional distribution on stock returns has revealed additional empirical regularities beyond simple volatility clustering. Negative return shocks are more likely than positive shocks to be followed by subsequent high volatility (see, for example, Bekaert and Wu 2000).

4 LARGE JUMPS

As mentioned above, jumps in the stock price will tend to cause fat tails in the distribution of returns. In addition, jumps have other important implications for risk management. The traditional approach to option pricing assumes that market participants may trade at any time, with no transaction costs. Coupled with the assumption of a continuous diffusion process, this allows the hypothetical arbitrageur to replicate an option using a continuously-updated trading strategy in the stock. Alternatively, in this fictional world, a risk manager could follow a "stop loss" strategy. For example, if you were to buy a stock at 100, you could guarantee that your loss on that position will never exceed ten percent by selling the stock the moment it reaches 90.

In practice, large jumps in prices can make it impossible to implement a dynamic replicating strategy or a stop-loss strategy. The stock may close above 90 one day and open below 90 the next day, with no opportunity to sell on the way down. Or, the price can jump down while the market is open.

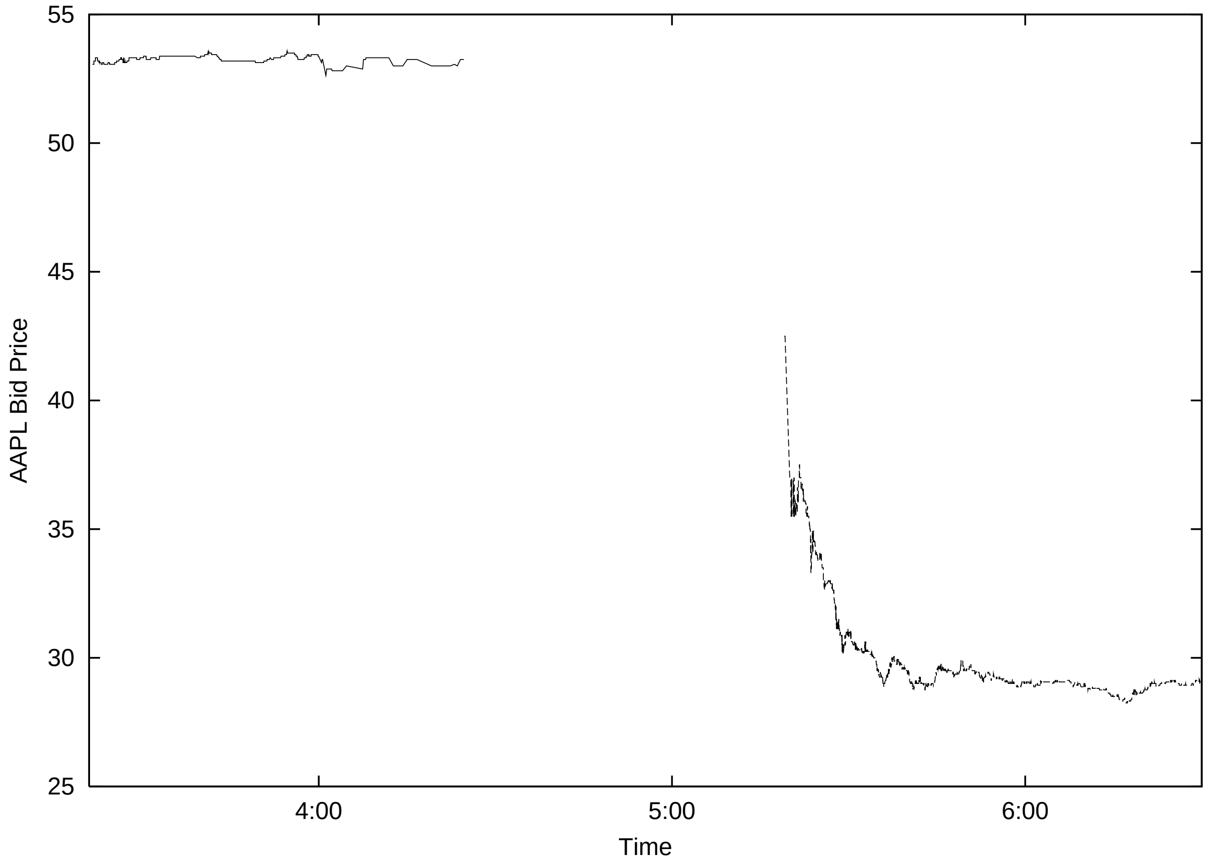

Figure 1: Large Jumps: Bid Prices for AAPL Stock, 9/28/2000

To illustrate, on September 28, 2000, after the close of regular trading hours, Apple Computer announced that their quarterly profits would be considerably lower than the market had anticipated. After-hours trading was immediately halted when the announcement hit the market, at around 4:24 P.M. When trading stopped, the bid price for Apple was at $53.25. When after-hours trading resumed at 5:18, the bid price immediately dropped to $42.50, and then less than sixty seconds later dropped to $37.00. By 6:30, the price had dropped to around $29. Such large sudden moves tend to wreak havoc on dynamic hedging or replicating strategies.

Another example involves the "Portfolio Insurance" strategy, which gained great popularity with institutional investors in the mid 1980s. The idea here was that institutions could protect their capital against a stock market crash, but still enjoy upside potential, by following a dynamic strategy that was essentially equivalent to buying a protective put on the portfolio. The strategy required the investor to move money out of cash into stocks as the market rose, and to sell the stocks as the market fell. During the crash of 1987, the strategy performed far worse than predicted, as investors found it extremely difficult to unload stocks as the price dropped. At times during the day, liquidity completely vanished, and investors could not sell at any price.

5 NON-NORMAL COPULAS

Up until now, our discussion has focused on properties that are readily apparent in univariate time series (fat tails, volatility clustering, jumps). But the heart of risk management is aggregation across several sources of risk, and this requires assumptions on the joint distribution of multiple state variables. For problems involving multiple sources of risk, the standard approach that has been used by practitioners and academics for many years is to estimate a variance-covariance (VC) matrix for the returns on all the securities in the portfolio. The VC matrix is an integral part of the machinery of modern portfolio theory, as developed by Markowitz (1952) and others.

If returns are assumed to be multivariate normal, the VC matrix is sufficient to identify the distribution, and thus can be used to calculate tail probabilities for a portfolio of positions in the individual securities. For more general distributions, this is no longer true. As we have seen, the univariate distributions are not normal, and recent history has taught us the dangers of using historical correlations to estimate VAR.

The dependence between two variables can be characterized more generally by the copula function—that is, the function relating the joint distribution to the marginal distribution. It is difficult to make blanket generalizations about joint distributions between financial time series, as there are many different combinations of variables, and only a few results have been published. With that caveat in mind, let us mention in passing a few results that deserve further attention.

There is ample evidence that the covariance between two financial variables is somewhat predictable. Just as the conditional variance is a function of past return shocks in univariate GARCH models, the conditional covariance in multivariate GARCH models is often related to recent shocks in both variables. Volatility shocks in the bond market can leak over to the stock market, and vice versa (Fleming, Kirby and Ostdiek 1998).

In addition, there is evidence of asymmetric dependence—the extent to which stock prices move together appears to be different for positive and negative price changes. Specifically, correlations between stocks seem to be higher when the market goes down. Also, prices tend to move together more closely on days of extreme market movements, particularly on the negative side. For more details, see Longin and Solnik (2001).

6 PARAMETER INSTABILITY

Markets are always changing. Regulatory structures evolve, new financial instruments are introduced, new trading strategies are invented, information technology advances, macroeconomic variables fluctuate, new industries come and go. It is not surprising that the parameters estimated from GARCH and other time-series models tend to fluctuate considerably over time (see, for example, Lamoureux and Lastrapes 1990). In the face of changing market conditions, the financial engineer must address the vital question of "how far back to look" when estimating volatility, correlation, or other parameters in time-series models.

The econometrician is stuck between Scylla and Charybdis. With too little data, the parameters cannot be estimated with any precision—in an idealized, unchanging world, five to ten years of daily returns may be required to obtain reliable estimates for GARCH and other more elaborate time-series models. With too much data, the model will be misspecified because the parameters are not stable over the time period. The best course depends on how quickly the world is changing. If there were one true and unchanging returns generating process, and we knew its form, we could get a more and more accurate estimate of its parameters by using more and more data. On the other extreme, the world may be changing so quickly that it is impossible to obtain reliable parameter estimates, and only the most recent data would be informative at all.

For measuring short-horizon volatility, the prevailing view among practitioners seems to be that only the very recent history is relevant. Thus, we see data providers like RiskMetrics providing volatility forecasts based on an exponentially-weighted moving average with a decay factor around .94, a procedure that places most of the weight on the most recent observations, and relatively little weight on events older than a few weeks.

It is not uncommon to see GARCH models estimated with samples as short as one or two years. Although long daily time series are usually available, the older observations are frequently ignored. Presumably, this is based on an underlying belief that market conditions change so quickly that there is little to learn from the older data. However, one should be cautious about using such short samples. In order to get a reliable measure of how quickly volatility shocks decay, one needs a sample long enough to contain a decent number of large shocks. The shorter the sample, the more noisy will be the estimates. Parameter estimates tend to be sensitive to the inclusion or exclusion of extreme events like the crash of 1987.

We must be especially wary of small samples when our goal is to estimate tail probabilities. How can one estimate the frequency of extreme events using a sample so small that it does not include any? One would not think of using only one month of data to estimate the distribution of earthquake sizes in San Francisco. For distributions with high skewness and/or kurtosis, it can be very difficult to obtain reliable estimates of these moments, even with a large number of observations. To highlight this problem, let us again return to our sample of 32,572 daily stock index returns from 1885 to 2002. Over this entire period, the sample skewness is approximately -0.133. If October 19, 1987 is omitted from the sample, the sample skewness becomes positive, at approximately 0.114. If we also omit the positive return on March 15, 1933, the estimate goes back down to -0.018. Even with well over a hundred years of returns data, our estimates of skewness are highly sensitive to the inclusion of one or two observations.

There is much to be learned from going back and examining a long history of stock returns. Most academic research on daily returns in the U.S. stock market uses the CRSP database, which begins in 1962. Restricting focus to the period since 1962 may have distorted our perceptions about markets in favor of characteristics that are most evident in that period.

Figure 2 depicts market volatility from 1885 to the present. On the Y axis is a rolling one-year estimate of the standard deviation of daily stock returns. The past few years have been quite volatile, one of the most volatile periods since World War II. Data from the great depression helps put these volatility numbers in perspective.

As another example, Ahn et al. (2002) state that "[a]rguably one of the most striking asset price anomalies is the evidence of large, positive short-horizon autocorrelations for returns on stock portfolios." Figure 3 depicts the sample

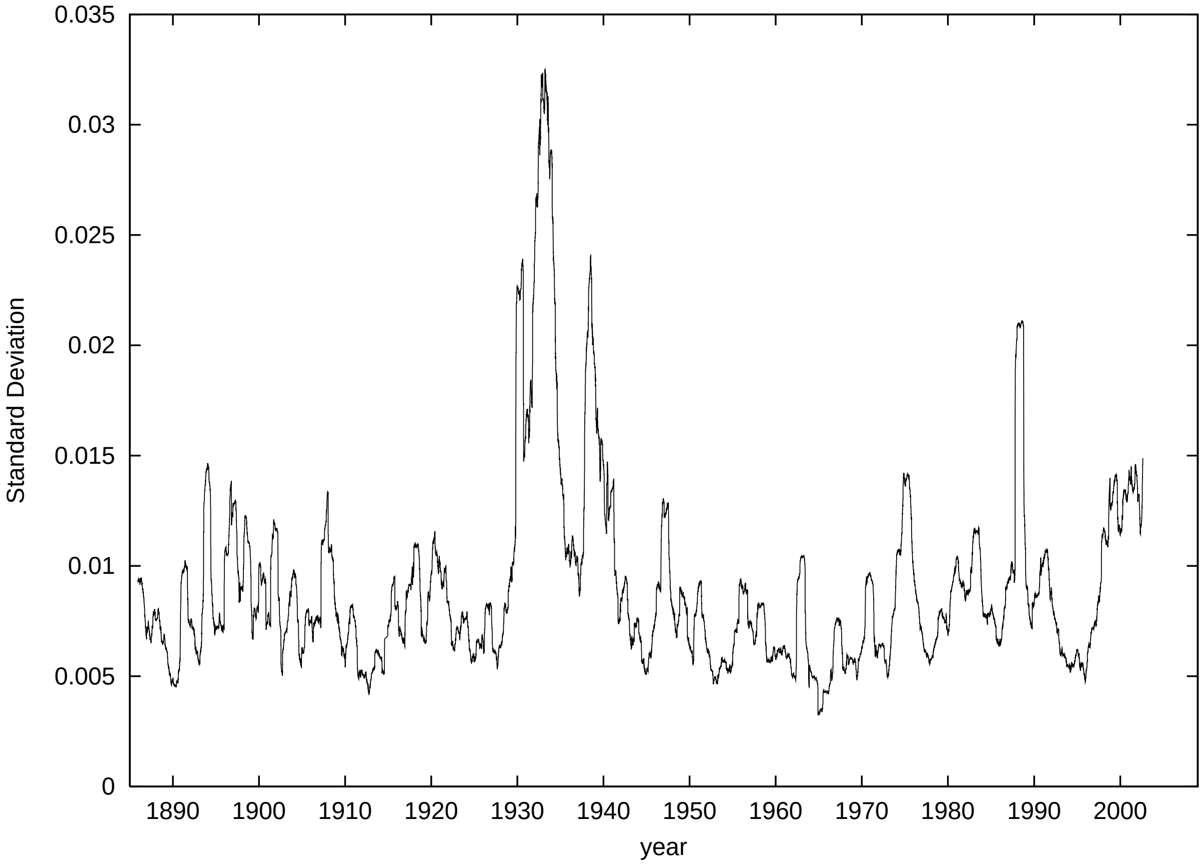

Figure 2: Standard Deviation of Daily Returns: Rolling One-Year Window

first-order autocorrelation for a moving window of 1,000 trading days since 1885. Certainly, positive autocorrelation appears to have been the norm since the end of World War II, and particularly in the 1970s. However, prior to the war there were several long periods when autocorrelation was near zero or negative. Interestingly, the most recent data seem to indicate that we may have moved back to a regime of autocorrelations near zero.

To summarize, regardless of what specific model is chosen, the parameters describing the dynamics of asset price returns tend to be hard to estimate, and they tend to be unstable over time. As a result, it is difficult to determine which of many competing models is most suitable for option pricing and risk management. No consensus has emerged as to which model is the most accurate, or about how much data should be used when estimating parameters. By testing our models over long historical time periods, we can observe the extent to which parameter estimates tend to vary over time, and attempt to ascertain whether these fluctuations correspond to true structural changes or to extreme events. Using simulation, we can explicitly measure the potential costs of model misspecification and parameter uncertainty, or generalize existing models to allow for randomly changing parameters.

REFERENCES

Ahn, D., J. Boudoukh, M. Richardson and R. F. Whitelaw. 2002. Partial Adjustment or Stale Prices? Implications from Stock Index and Futures Return Autocorrelations. *Review of Financial Studies* 16: 655–689.

Bakshi, G., C. Cao and Z. Chen. 1997. Empirical Performance of Alternative Option Pricing Models. *Journal of Finance* 52: 2003–2049.

Bekaert, G. and G. Wu. 2000. Asymmetric Volatility and Risk in Equity Markets. *Review of Financial Studies* 13: 1–42.

Bollerslev, T. 1986. Generalized Autoregressive Conditional Heteroskedasticity. *Journal of Econometrics* 31: 307–327.

Engle, R.F. 1982. Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom Inflation, *Econometrica* 50: 987–1007.

Fama, E.F. 1965. The Behavior of Stock Market Prices. *Journal of Business* 38: 34–105.

Fleming, J., C. Kirby, and B. Ostdiek. 1998. Information and Volatility Linkages in the Stock, Bond, and Money Markets. *Journal of Financial Economics* 49: 111–137.

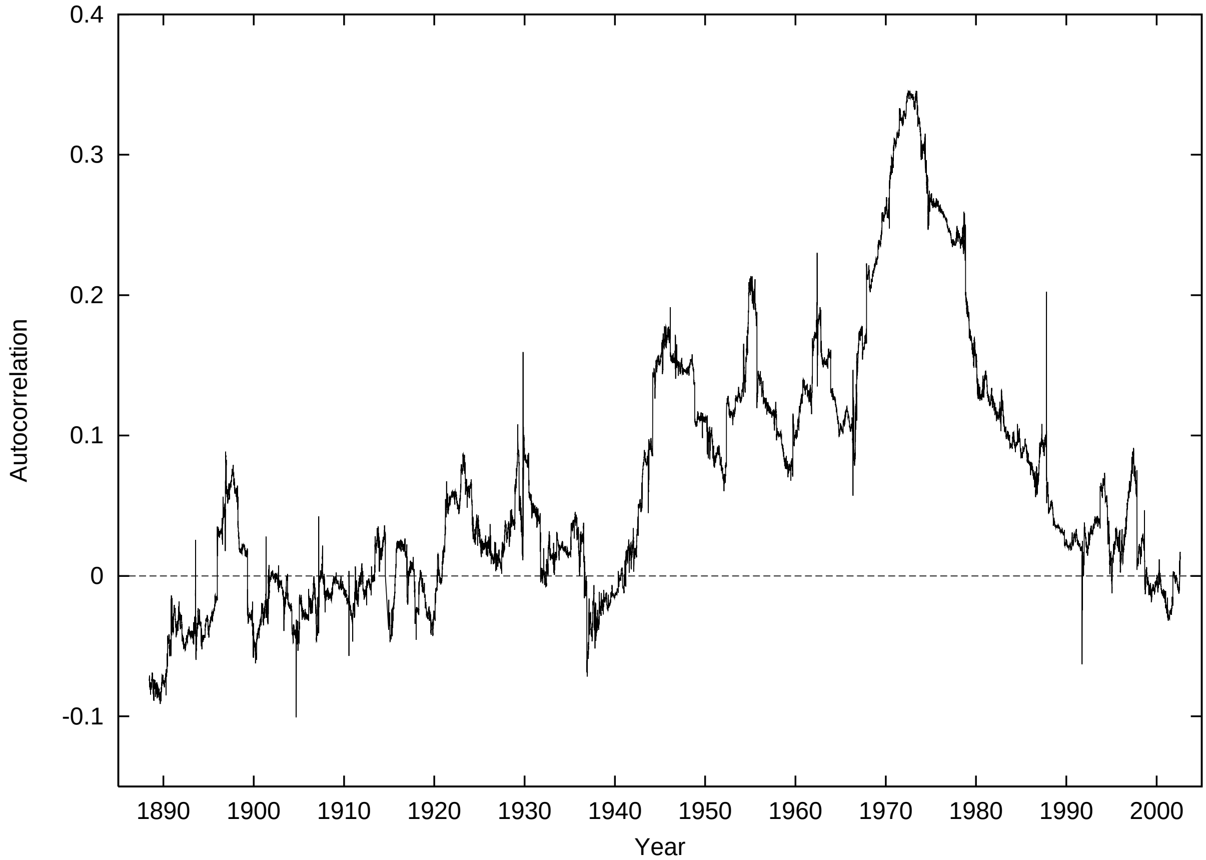

Figure 3: First-order Autocorrelation in Daily Returns: Rolling 1000-Day Window

Hamilton, J. 1994. *Time Series Analysis*, Princeton University Press.

Heston, S. 1993. A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bond Currency Options. *Review of Financial Studies* 6: 327–343.

Lamoureux, C. G. and W. D. Lastrapes. 1990. Persistence in Variance, Structural Change, and the GARCH Model. *Journal of Business and Economic Statistics* 68:225-234.

Longin, F. and B. Solnik. 2001. Extreme Correlation of International Equity Markets. *Journal of Finance* 56: 651–678.

Lowenstein, R. 2000. *When Genius Failed: The Rise and Fall of Long-Term Capital Management*. Random House. New York.

Mandelbrot B. 1963. The Variation of Certain Speculative Prices. *Journal of Business* 36: 394–419.

Markowitz, H. 1952. Portfolio Selection. *Journal of Finance* 17: 77–91.

Schwert, G.W. 1990. Indexes of U.S. Stock Prices from 1802 to 1987. *Journal of Business* 63: 399–426.

AUTHOR BIOGRAPHY

STEWART MAYHEW is an Assistant Professor of Finance in the Terry College of Business at the University of Georgia. His e-mail address is <mayhew@terry.uga.edu>.