J.P.Morgan/Reuters

RiskMetrics™—Technical Document

Fourth Edition, 1996

New York
December 17, 1996

- J.P. Morgan and Reuters have teamed up to enhance RiskMetrics™. Morgan will continue to be responsible for enhancing the methods outlined in this document, while Reuters will control the production and distribution of the RiskMetrics™ data sets.
- Expanded sections on methodology outline enhanced analytical solutions for dealing with nonlinear options risks and introduce methods on how to account for non-normal distributions.
- Enclosed diskette contains many examples used in this document. It allows readers to experiment with our risk measurement techniques.
- All publications and daily data sets are available free of charge on J.P. Morgan's Web page on the Internet at **http://www.jpmorgan.com/RiskManagement/RiskMetrics/RiskMetrics.html**. This page is accessible directly or through third party services such as CompuServe®, America Online™, or Prodigy®.

Morgan Guaranty Trust Company
Risk Management Advisory
Jacques Longerstaey
(1-212) 648-4936
riskmetrics@jpmorgan.com

Reuters Ltd
International Marketing
Martin Spencer
(44-171) 542-3260
martin.spencer@reuters.com

This *Technical Document* provides a detailed description of RiskMetrics™, a set of techniques and data to measure market risks in portfolios of fixed income instruments, equities, foreign exchange, commodities, and their derivatives issued in over 30 countries. This edition has been expanded significantly from the previous release issued in May 1995.

We make this methodology and the corresponding RiskMetrics™ data sets available for three reasons:

1. We are interested in promoting greater transparency of market risks. Transparency is the key to effective risk management.

2. Our aim has been to establish a benchmark for market risk measurement. The absence of a common point of reference for market risks makes it difficult to compare different approaches to and measures of market risks. Risks are comparable only when they are measured with the same yardstick.

3. We intend to provide our clients with sound advice, including advice on managing their market risks. We describe the RiskMetrics™ methodology as an aid to clients in understanding and evaluating that advice.

Both J.P. Morgan and Reuters are committed to further the development of RiskMetrics™ as a fully transparent set of risk measurement methods. We look forward to continued feedback on how to maintain the quality that has made RiskMetrics™ the benchmark for measuring market risk.

RiskMetrics™ is based on, but differs significantly from, the risk measurement methodology developed by J.P. Morgan for the measurement, management, and control of market risks in its trading, arbitrage, and own investment account activities. **We remind our readers that no amount of sophisticated analytics will replace experience and professional judgment in managing risks**. RiskMetrics™ is nothing more than a high-quality tool for the professional risk manager involved in the financial markets and is not a guarantee of specific results.

RiskMetrics™—Technical Document
Fourth Edition (December 1996)

This book

This is the reference document for RiskMetrics™. It covers all aspects of RiskMetrics and supersedes all previous editions of the *Technical Document*. It is meant to serve as a reference to the methodology of statistical estimation of market risk, as well as detailed documentation of the analytics that generate the data sets that are published daily on our Internet Web sites.

This document reviews

1. The conceptual framework underlying the methodologies for estimating market risks.
2. The statistics of financial market returns.
3. How to model financial instrument exposures to a variety of market risk factors.
4. The data sets of statistical measures that we estimate and distribute daily over the Internet and shortly, the Reuters Web.

Measurement and management of market risks continues to be as much a craft as it is a science. It has evolved rapidly over the last 15 years and has continued to evolve since we launched RiskMetrics in October 1994. Dozens of professionals at J.P. Morgan have contributed to the development of this market risk management technology and the latest document contains entries or contributions from a significant number of our market risk professionals.

We have received numerous constructive comments and criticisms from professionals at Central Banks and regulatory bodies in many countries, from our competitors at other financial institutions, from a large number specialists in academia and last, but not least, from our clients. Without their feedback, help, and encouragement to pursue our strategy of open disclosure of methodology and free access to data, we would not have been as successful in advancing this technology as much as we have over the last two years.

What is RiskMetrics?

RiskMetrics is a set of tools that enable participants in the financial markets to estimate their exposure to market risk under what has been called the “Value-at-Risk framework”. RiskMetrics has three basic components:

- A set of market risk measurement methodologies outlined in this document.
- Data sets of volatility and correlation data used in the computation of market risk.
- Software systems developed by J.P.Morgan, subsidiaries of Reuters, and third party vendors that implement the methodologies described herein.

With the help of this document and the associated line of products, users should be in a position to estimate market risks in portfolios of foreign exchange, fixed income, equity and commodity products.

J.P. Morgan and Reuters team up on RiskMetrics

In June 1996, J.P. Morgan signed an agreement with Reuters to cooperate on the building of a new and more powerful version of RiskMetrics. Since the launch of RiskMetrics in October 1994, we have received numerous requests to add new products, instruments, and markets to the daily volatility and correlation data sets. We have also perceived the need in the market for a more flexible VaR data tool than the standard matrices that are currently distributed over the Internet. The new partnership with Reuters, which will be based on the precept that both firms will focus on their respective strengths, will help us achieve these objectives.

Methodology

J.P. Morgan will continue to develop the RiskMetrics set of VaR methodologies and publish them in the quarterly *RiskMetrics Monitor* and in the annual *RiskMetrics—Technical Document*.

RiskMetrics data sets

Reuters will take over the responsibility for data sourcing as well as production and delivery of the risk data sets. The current RiskMetrics data sets will continue to be available on the Internet free of charge and will be further improved as a benchmark tool designed to broaden the understanding of the principles of market risk measurement.

When J.P. Morgan first launched RiskMetrics in October 1994, the objective was to go for broad market coverage initially, and follow up with more granularity in terms of the markets and instruments covered. This over time, would reduce the need for proxies and would provide additional data to measure more accurately the risk associated with non-linear instruments.

The partnership will address these new markets and products and will also introduce a new customizable service, which will be available over the Reuters Web service. The customizable RiskMetrics approach will give risk managers the ability to scale data to meet the needs of their individual trading profiles. Its capabilities will range from providing customized covariance matrices needed to run VaR calculations, to supplying data for historical simulation and stress-testing scenarios.

More details on these plans will be discussed in later editions of the *RiskMetrics Monitor.*

Systems

Both J.P. Morgan and Reuters, through its Sailfish subsidiary, have developed client-site RiskMetrics VaR applications. These products, together with the expanding suite of third party applications will continue to provide RiskMetrics implementations.

What is new in this fourth edition?

In terms of content, the Fourth Edition of the *Technical Document* incorporates the changes and refinements to the methodology that were initially outlined in the 1995–1996 editions of the *RiskMetrics Monitor*:

- **Expanded framework:** We have worked extensively on refining the analytical framework for analyzing options risk without having to perform relatively time consuming simulations and have outlined the basis for an improved methodology which incorporates better information on the tails of distributions related to financial asset price returns; we've also developed a data synchronization algorithm to refine our volatility and correlation estimates for products which do not trade in the same time zone;

- **New markets:** We expanded the daily data sets to include estimated volatilities and correlations of additional foreign exchange, fixed income and equity markets, particularly in South East Asia and Latin America.

- **Fine-tuned methodology:** We have modified the approach in a number of ways. First, we've changed our definition of price volatility which is now based on a total return concept; we've also revised some of the algorithms used in our mapping routines and are in the process of redefining the techniques used in estimating equity portfolio risk.

- **RiskMetrics products:** While we have continued to expand the list of third parties providing RiskMetrics products and support, this is no longer included with this document. Given the rapid pace of change in the availability of risk management software products, readers are advised to consult our Internet web site for the latest available list of products. This list, which now includes FourFifteen™, J.P. Morgan's own VaR calculator and report generating software, continues to grow, attesting to the broad acceptance RiskMetrics has achieved.

- **New tools to use the RiskMetrics data sets:** We have published an Excel add-in function which enables users to import volatilities and correlations directly into a spreadsheet. This tool is available from our Internet web site.

The structure of the document has changed only slightly. As before, its size warrants the following note: One need not read and understand the entire document in order to benefit from RiskMetrics. The document is organized in parts that address subjects of particular interest to many readers.

Part I: Risk Measurement Framework

This part is for the general practitioner. It provides a practical framework on how to think about market risks, how to apply that thinking in practice, and how to interpret the results. It reviews the different approaches to risk estimation, shows how the calculations work on simple examples and discusses how the results can be used in limit management, performance evaluation, and capital allocation.

Part II: Statistics of Financial Market Returns

This part requires an understanding and interest in statistical analysis. It reviews the assumptions behind the statistics used to describe financial market returns and how distributions of future returns can be estimated.

Part III: Risk Modeling of Financial Instruments

This part is required reading for implementation of a market risk measurement system. It reviews how positions in any asset class can be described in a standardized fashion (foreign exchange, interest rates, equities, and commodities). Special attention is given to derivatives positions. The purpose is to demystify derivatives in order to show that their market risks can be measured in the same fashion as their underlying.

Part IV: RiskMetrics Data Sets

This part should be of interest to users of the RiskMetrics data sets. First it describes the sources of all daily price and rate data. It then discusses the attributes of each volatility and correlation series in the RiskMetrics data sets. And last, it provides detailed format descriptions required to decipher the data sets that can be downloaded from public or commercial sources.

Appendices

This part reviews some of the more technical issues surrounding methodology and regulatory requirements for market risk capital in banks and demonstrates the use of RiskMetrics with the example diskette provided with this document. Finally, Appendix H shows you how to access the RiskMetrics data sets from the Internet.

RiskMetrics examples diskette

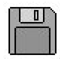

This diskette is located inside the back cover. It contains an Excel workbook that includes some of the examples shown in this document. Such examples are identified by the icon shown here.

Future plans

We expect to update this *Technical Document* annually as we adapt our market risk standards to further improve the techniques and data to meet the changing needs of our clients.

RiskMetrics is a now an integral part of J.P. Morgan's Risk Management Services group which provides advisory services to a wide variety of the firm's clients. We continue to welcome any suggestions to enhance the methodology and adapt it further to the needs of the market. All suggestions, requests and inquiries should be directed to the authors of this publication or to your local RiskMetrics contacts listed on the back cover.

Acknowledgments

The authors would like to thank the numerous individuals who participated in the writing and editing of this document, particularly Chris Finger and Chris Athaide from J.P. Morgan's risk management research group, and Elizabeth Frederick and John Matero from our risk advisory practice. Finally, this document could not have been produced without the contributions of our consulting editor, Tatiana Kolubayev. We apologize for any omissions to this list.

Table of contents

Part IV RiskMetrics Data Sets

Part V Backtesting

Appendices

Reference

List of charts

List of tables

Part II

Statistics of Financial Market Returns

Chapter 4. Statistical and probability foundations

Chapter 4. Statistical and probability foundations

Peter Zangari
Morgan Guaranty Trust Company
Risk Management Research
(1-212) 648-8641
zangari_peter@jpmorgan.com

This chapter presents the statistical and probability underpinnings of the RiskMetrics model. It explains the assumptions commonly applied to forecast the distribution of portfolio returns and investigates the empirical validity of these assumptions. While we have tried to make this chapter self-contained, its subject matter does require a thorough grasp of elementary statistics. We have included many up-to-date references on specific topics so that the interested reader may pursue further study in these areas.

This chapter is organized as follows:

- Section 4.1 presents definitions of financial price returns and explains the type of returns applied in RiskMetrics.
- Section 4.2 describes the basic random walk model for financial prices to serve as background to introducing the RiskMetrics model of returns.
- Section 4.3 looks at some observed time series properties of financial returns in the context of the random walk model.
- Section 4.4 summarizes the results presented in Sections 4.1 through 4.3.
- Section 4.5 reviews some popular models of financial returns and presents a review of the normal and lognormal distributions.
- Section 4.6 presents the RiskMetrics model as a modified random walk. This section lists the assumptions of the RiskMetrics model—that is, what RiskMetrics assumes about the evolution of financial returns over time and the distribution of returns at any point in time.
- Section 4.7 is a chapter summary.

4.1 Definition of financial price changes and returns[1]

Risk is often measured in terms of price changes. These changes can take a variety of forms such as absolute price change, relative price change, and log price change. When a price change is defined relative to some initial price, it is known as a return. **RiskMetrics measures change in value of a portfolio (often referred to as the adverse price move) in terms of log price changes also known as continuously-compounded returns.** Next, we explain different definitions of price returns.

4.1.1 One-day (single period) horizon

Denote by P_t the price of a security at date *t*. In this document, *t* is taken to represent one business day.

The absolute price change on a security between dates *t* and $t-1$ (i.e., one day) is defined as

[4.1] $$D_t = P_t - P_{t-1}$$

[1] References for this section are, Campbell, Lo and MacKinley (1995) and Taylor, S. J. (1987).

The relative price change, or percent return[2], R_t, for the same period is

$$R_t = \frac{P_t - P_{t-1}}{P_{t-1}} \quad [4.2]$$

If the gross return on a security is just $1 + R_t$, then the log price change (or continuously-compounded return), r_t, of a security is defined to be the natural logarithm of its gross return. That is,

$$\begin{aligned} r_t &= \ln(1 + R_t) \\ &= \ln\left(\frac{P_t}{P_{t-1}}\right) \\ &= (p_t - p_{t-1}) \end{aligned} \quad [4.3]$$

where $p_t = \ln(P_t)$ is the natural logarithm of P_t.

In practice, the main reason for working with returns rather than prices is that returns have more attractive statistical properties than prices, as will be shown below. Further, returns (relative and log price changes) are often preferred to absolute price changes because the latter do not measure change in terms of the **given** price level.

To illustrate the different results that different price changes can yield, Table 4.1 presents daily USD/DEM exchange rates for the period 28-Mar-96 through 12-Apr-96 and the corresponding daily absolute, relative, and log price changes.

Table 4.1
Absolute, relative and log price changes*

Date	Price (USD/DEM), P_t	Absolute price change (%), D_t	Relative price change (%), R_t	Log price change (%,) r_t
28-Mar-96	0.67654	0.427	0.635	0.633
29-Mar-96	0.67732	0.078	0.115	0.115
1-Apr-96	0.67422	–0.310	–0.458	–0.459
2-Apr-96	0.67485	0.063	0.093	0.093
3-Apr-96	0.67604	0.119	0.176	0.176
4-Apr-96	0.67545	–0.059	–0.087	–0.087
5-Apr-96	0.67449	–0.096	–0.142	-0.142
8-Apr-96	0.67668	0.219	0.325	0.324
9-Apr-96	0.67033	–0.635	–0.938	–0.943
10-Apr-96	0.66680	–0.353	–0.527	–0.528
11-Apr-96	0.66609	–0.071	–0.106	–0.107
12-Apr-96	0.66503	–0.106	–0.159	–0.159

* RiskMetrics foreign exchange series are quoted as USD per unit foreign currency given that the datasets are standardized for users whose base currency is the USD. This is the inverse of market quotation standards for most currency pairs.

As expected, all three series of price changes have the same sign for any given day. Also, notice the similarity between the log and relative price changes. In fact, we should expect these two return series to be similar to one another for small changes in the underlying prices. In contrast, the absolute change series is quite different from the other two series.

[2] Although it is called "percent return," the relative price change is expressed as a decimal number.

To further illustrate the potential differences between absolute and log price changes, Chart 4.1 shows daily absolute and log price changes for the U.S. 30-year government bond over the first quarter of 1996.

Chart 4.1
Absolute price change and log price change in U.S. 30-year government bond

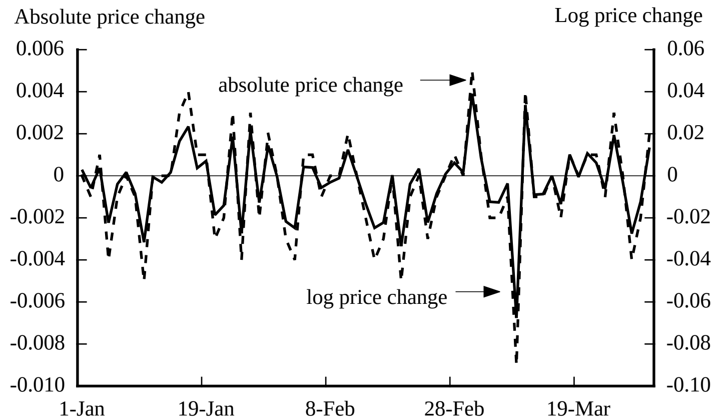

Chart 4.1 shows that movements of the two changes over time are quite similar although the magnitude of their variation is different. This latter point and the results presented in Table 4.1 should make it clear that it is important to understand the convention chosen for measuring price changes.

4.1.2 Multiple-day (multi-period) horizon

The returns R_t and r_t described above are 1-day returns. We now show how to use them to compute returns for horizons greater than one day.

Multiple-day percent returns over the most recent *k* days, $R_t(k)$, are defined simply as

[4.4] $$R_t(k) = \frac{P_t - P_{t-k}}{P_{t-k}}$$

In terms of 1-day returns, the multiple-day **gross** return $1 + R_t(k)$ is given by the product of 1-day gross returns.

[4.5] $$\begin{aligned} 1 + R_t(k) &= (1 + R_t)(1 + R_{t-1}) \cdots (1 + R_{t-k-1}) \\ &= \frac{P_t}{P_{t-1}} \cdot \frac{P_{t-1}}{P_{t-2}} \cdot \ldots \cdot \frac{P_{t-k-1}}{P_{t-k}} \\ &= \frac{P_t}{P_{t-k}} \end{aligned}$$

Note that in Eq. [4.5] the *k*-day return is a discretely compounded return. For continuously compounded returns, the multiple-day return $r_t(k)$ is defined as

[4.6] $$r_t(k) = \ln\left(\frac{P_t}{P_{t-k}}\right)$$

The continuously-compounded return $r_t(k)$ is the sum of k continuously-compounded 1-day returns. To see this we use the relation $r_t(k) = \ln[1 + R_t(k)]$. The return $r_t(k)$ can then be written as

[4.7]
$$\begin{aligned} r_t(k) &= \ln[1 + R_t(k)] \\ &= \ln[(1 + R_t) \cdot (1 + R_{t-1}) \cdot (1 + R_{t-k-1})] \\ &= r_t + r_{t-1} + \ldots + r_{t-k+1} \end{aligned}$$

Notice from Eq. [4.7] that compounding, a multiplicative operation, is converted to an additive operation by taking logarithms. Therefore, multiple day returns based on continuous compounding are simple sums of one-day returns.

As an example of how 1-day returns are used to generate a multiple-day return, we use a 1-month period, defined by RiskMetrics as having 25 business days. Working with log price changes, the continuously compounded return over one month is given by

[4.8]
$$r_t(25) = r_t + r_{t-1} + \ldots + r_{t-24}$$

That is, the 1-month return is the sum of the last 25 1-day returns.

4.1.3 Percent and continuous compounding in aggregating returns

When deciding whether to work with percent or continuously compounded returns it is important to understand how such returns aggregate both across time and across individual returns at any point in time.

In the preceding section we showed how multiple-day returns can be constructed from 1-day returns by aggregating the latter across time. This is known as temporal aggregation. However, there is another type of aggregation known as cross-section aggregation. In the latter approach, aggregation is across individual returns (each corresponding to a specific instrument) at a particular point in time. For example, consider a portfolio that consists of three instruments. Let r_i and R_i $(i = 1, 2, 3)$ be the continuously compounded and percent returns, respectively and let w_i represent the portfolio weights. (The parameter w_i represents the fraction of the total portfolio value allocated to the *i*th instrument with the condition that—assuming no short positions— $w_1 + w_2 + w_3 = 1$). If the initial value of this portfolio is P_0 the price of the portfolio one period later with continuously compounded returns is

[4.9]
$$P_1 = w_1 \cdot P_0 \cdot e^{r_1} + w_2 \cdot P_0 \cdot e^{r_2} + w_3 \cdot P_0 \cdot e^{r_3}$$

Solving Eq. [4.9] for the portfolio return, $r_p = \ln\left(\frac{P_1}{P_0}\right)$, we get

[4.10]
$$r_p = \ln\left(w_1 \cdot e^{r_1} + w_2 \cdot e^{r_2} + w_3 \cdot e^{r_3}\right)$$

The price of the portfolio one period later with discrete compounding, i.e., using percent returns, is

[4.11]
$$P_1 = w_1 \cdot P_0 \cdot (1 + r_1) + w_2 \cdot P_0 \cdot (1 + r_2) + w_3 \cdot P_0 \cdot (1 + r_3)$$

The percent portfolio return, $R_p = \frac{(P_1 - P_0)}{P_0}$, is given by

[4.12]
$$R_p = w_1 \cdot r_1 + w_2 \cdot r_2 + w_3 \cdot r_3$$

Equation [4.12] is the expression often used to describe a portfolio return—as a weighted sum of individual returns.

Table 4.2 presents expressions for returns that are constructed from temporal and cross-section aggregation for percent and continuously compounded returns.

Table 4.2
Return aggregation

Aggregation	Temporal	Cross-section
Percent returns	$R_{it}(k) = \prod_{t=1}^{T} (1 + R_{it}) - 1$	$R_{pt} = \sum_{i=1}^{N} w_i R_{it}$
Continuously compounded returns	$r_{it}(k) = \sum_{t=1}^{T} r_{it}$	$r_{pt} = \ln\left(\sum_{i=1}^{N} w_i e^{r_{it}}\right)$

The table shows that when aggregation is done across time, it is more convenient to work with continuously compounded returns whereas when aggregation is across assets, percent returns offer a simpler expression.

As previously stated, log price changes (continuously compounded returns) are used in RiskMetrics as the basis for all computations. In practice, RiskMetrics assumes that a portfolio return is a weighted average of continuously compounded returns. That is, a portfolio return is defined as follows

[4.13] $$r_{pt} \cong \sum_{i=1}^{N} w_i r_{it}$$

As will be discussed in detail in the next section, when 1-day returns are computed using r_t, then a model describing the distribution of 1-day returns extends straightforwardly to returns greater than one day.[3]

In the next two sections (4.2 and 4.3) we describe a class of time series models and investigate the empirical properties of financial returns. These sections serve as important background to understanding the assumptions RiskMetrics applies to financial returns.

4.2 Modeling financial prices and returns

A risk measurement model attempts to characterize the future change in a portfolio's value. Often, it does so by making forecasts of each of a portfolio's underlying instrument's future price changes, using only past changes to construct these forecasts. This task of describing future price changes requires that we model the following; (1) the temporal dynamics of returns, i.e., model the evolution of returns over time, and (2) the distribution of returns at any point in time.

A widely used class of models that describes the evolution of price returns is based on the notion that financial prices follow a random walk.

[3] There are two other reasons for using log price changes. The first relates to "Siegel's paradox," Meese, R.A. and Rogoff, K. (1983). The second relates to preserving normality for FX cross rates. Simply put, when using log price changes, FX cross rates can be written as differences of base currency rates. (See Section 8.4 for details.)

4.2.1 Random walk model for single-price assets

In this section we present a model for a security with a single price. Such a model applies naturally to assets such as foreign exchange rates, commodities, and equities where only one price exists per asset. The fundamental model of asset price dynamics is the random walk model,

$$\text{[4.14]}\qquad \begin{aligned} P_t &= \mu + P_{t-1} + \sigma\varepsilon_t \\ P_t - P_{t-1} &= \mu + \sigma\varepsilon_t, \qquad \varepsilon_t \sim \text{IID } N(0, 1) \end{aligned}$$

where IID stands for "identically and independently distributed"[4], and $N(0, 1)$ stands for the normal distribution with mean 0 and variance 1. Eq. [4.14] posits the evolution of prices and their distribution by noting that at any point in time, the current price P_t depends on a fixed parameter μ, last period's price P_{t-1}, and a normally distributed random variable, ε_t. Simply put, μ and σ affect the mean and variance of P_t's distribution, respectively.

The conditional distribution of P_t, given P_{t-1}, is normally distributed.[5] An obvious drawback of this model is that there will always be a non-zero probability that prices are negative.[6] One way to guarantee that prices will be non-negative is to model the log price p_t as a random walk with normally distributed changes.

$$\text{[4.15]}\qquad p_t = \mu + p_{t-1} + \sigma\varepsilon_t \qquad \varepsilon_t \sim \text{IID } N(0, 1)$$

Notice that since we are modeling log prices, Eq. [4.15] is a model for continuously compounded returns, i.e., $r_t = \mu + \sigma\varepsilon_t$. Now, we can derive an expression for prices, P_t given last period's price P_{t-1} from Eq. [4.15]:

$$\text{[4.16]}\qquad P_t = P_{t-1}\exp(\mu + \sigma\varepsilon_t)$$

where $\exp(x) \equiv e^x$ and $e \cong 2.718$.

Since both P_{t-1} and $\exp(\mu + \sigma\varepsilon_t)$ are non-negative, we are guaranteed that P_t will never be negative. Also, when ε_t is normally distributed, P_t follows a lognormal distribution.[7]

Notice that both versions of the random walk model above assume that the change in (log) prices has a constant variance (i.e., σ does not change with time). We can relax this (unrealistic) assumption, thus allowing the variance of price changes to vary with time. Further, the variance could be modeled as a function of past information such as past variances. By allowing the variance to vary over time we have the model

$$\text{[4.17]}\qquad p_t = \mu + p_{t-1} + \sigma_t\varepsilon_t \qquad \varepsilon_t \sim N(0, 1)$$

---

[4] See Section 4.3 for the meaning of these assumptions.

[5] The unconditional distribution of P_t is undefined in that its mean and variance are infinite. This can easily be seen by solving Eq. [4.14] for P_t as a function of past ε_t's.

[6] This is because the normal distribution places a positive probability on all points from negative to positive infinity. See Section 4.5.2 for a discussion of the normal distribution.

[7] See Section 4.5.3 for a complete description of the lognormal distribution.

This version of the random walk model is important since it will be shown below that **RiskMetrics assumes that log prices evolve according to Eq. [4.17] with the parameter** μ **set to zero**.

4.2.2 Random walk model for fixed income instruments

With fixed income instruments we observe both prices and yields. When prices and yields exist, we must decide whether to model the log changes in the yields or in the prices. For example, for bonds, a well documented shortcoming of modeling price returns according to Eq. [4.15] is that the method ignores a bond's price **pull to par phenomenon**. That is, a bond has the distinct feature that as it approaches maturity, its price converges to its face value. Consequently, the bond price volatility will converge to zero.

Therefore, when modeling the dynamic behavior of bonds (and other fixed income instruments), the bond yields rather than the bond prices are often modeled according to the lognormal distribution. That is, if Y_t denotes the yield on a bond at period t, then $y_t = \ln(Y_t)$ is modeled as

[4.18] $$y_t = \mu + y_{t-1} + \sigma\varepsilon_t \qquad \varepsilon_t \sim IID \;\; \mathrm{N}(0, 1)$$

(Note that similar to Eq. [4.17] we can incorporate a time-varying variance into Eq. [4.18]). In addition to accounting for the pull to par phenomenon, another important reason for modeling the yield rather than the price according to Eq. [4.18] is that positive yields are guaranteed. In the context of bond option pricing, a strong case can often be made for modeling yields as lognormal.[8]

4.2.3 Time-dependent properties of the random walk model

Each of the random walk models presented in Sections 4.2.1 and 4.2.2 imply a certain movement in financial prices over time. In this section we use Eq. [4.15]—the random walk model in log prices, p_t—to explain some important properties of price dynamics implied by the random walk model. Specifically, we discuss the properties of stationary (mean-reverting) and nonstationary time series.

A stationary process is one where the mean and variance are constant and finite over time.[9] In order to introduce the properties of a stationary time series we must first generalize Eq. [4.15] to the following model.

[4.19] $$p_t = \mu + c \cdot p_{t-1} + \varepsilon_t \qquad \varepsilon_t \sim IID \;\; \mathrm{N}(0, 1),\; p_0 = 0$$

where c is a parameter. Here, a stationary time series is generated when $-1 < c < 1$. For example, if we set $c = 0.5$, we can simulate a stationary time series using

[4.20] $$p_t = 0.01 + 0.5 p_{t-1} + \varepsilon_t \qquad \varepsilon_t \sim IID \;\; \mathrm{N}(0, 1),\; p_0 = 0$$

[8] For a discussion on the potential advantages of modeling yield levels as lognormal, see Fabozzi (1989, Chapter 3).

[9] Stationarity also requires that the (auto-)covariance of returns at different times is only a function of the time between the returns, and not the times at which they occur. This definition of stationarity is known as weak or covariance stationarity.

Chart 4.2 shows the simulated stationary time series based on 500 simulations.

Chart 4.2
Simulated stationary/mean-reverting time series

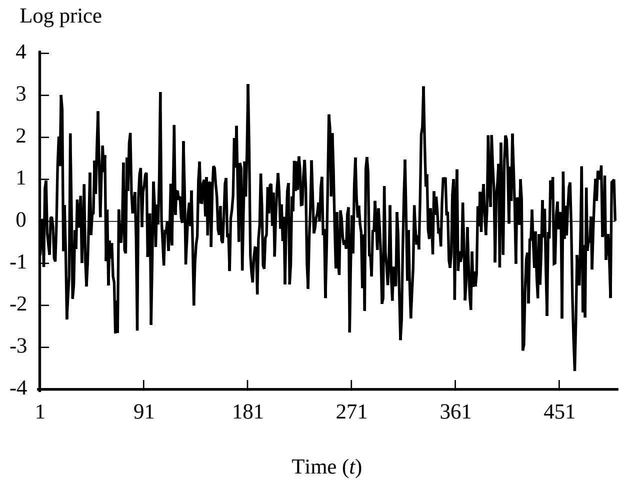

Chart 4.2 shows how a stationary series fluctuates around its mean, which in this model is 0.02. Hence, stationary series are **mean-reverting** since, regardless of the fluctuations' amplitudes, the series reverts to its mean.

Unlike a mean-reverting time series, a nonstationary time series does not fluctuate around a fixed mean. For example, in Eq. [4.15] the mean and variance of the log price p_t conditional on some original observed price, say p_0, are given by the following expressions

[4.21]
$$E_0[p_t|p_0] = p_0 + \mu t \qquad \text{(mean)}$$
$$V_0[p_t|p_0] = \sigma^2 t \qquad \text{(variance)}$$

where $E_0[\ ]$ and $V_0[\ ]$ are the expectation and variance operators taken at time 0. Eq. [4.21] shows that both the mean and variance of the log price are a function of time such that, as time t increases, so does p_t's conditional mean and variance. The fact that its mean and variance change with time and "blow-up" as time increases is a characteristic of a nonstationary time series.

To illustrate the properties of a nonstationary time series, we use the random walk model, Eq. [4.15], to simulate 500 data points. Specifically, we simulate a series based on the following model,

[4.22]
$$p_t = 0.01 + p_{t-1} + \varepsilon_t \qquad \varepsilon_t \sim IID\ \mathrm{N}(0,1),\ p_0 = 0$$

The simulated series is shown in Chart 4.3.

Chart 4.3
Simulated nonstationary time series

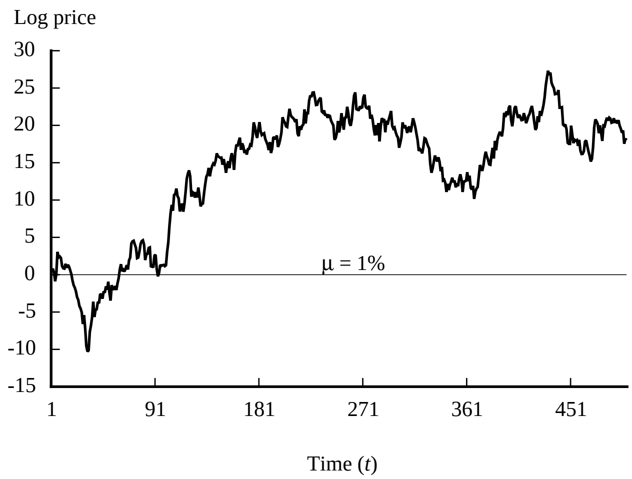

Notice how this series has a positive drift that grows with time, representing the term μt in Eq. [4.21]. This is a typical feature of a nonstationary time series.

In the preceding examples, notice that the difference between these stationary and nonstationary series is driven by the coefficient on last period's log price p_{t-1}. When this coefficient is 1, as in Eq. [4.22], the process generating log prices is known to have a "unit root". As should be expected, given the differences between stationary and non-stationary times series and their implications for statistical analysis, there is a large body of literature devoted to testing for the presence of a unit root.[10]

Real world examples of stationary and nonstationary series are shown in Charts 4.4 and 4.5. For the same period, Chart 4.4 plots the USD 30-year rate, a stationary time series.

Chart 4.4
Observed stationary time series
USD 30-year yield

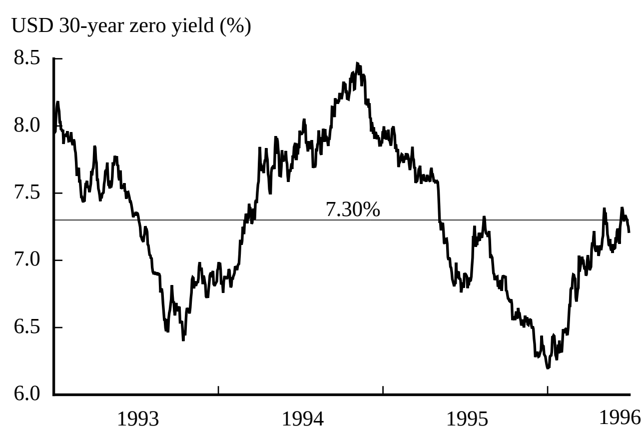

[10] A common statistical test for a unit root is known as the augmented Dickey-Fuller test. See Greene, (1993).

Notice how the 30-year rates fluctuate around the sample average of 7.30%, signifying that the time series for this period is mean-reverting.

Chart 4.5 plots the S&P 500 index for the period January 4, 1993 through June 28, 1996.

Chart 4.5
Observed nonstationary time series
S&P 500 index

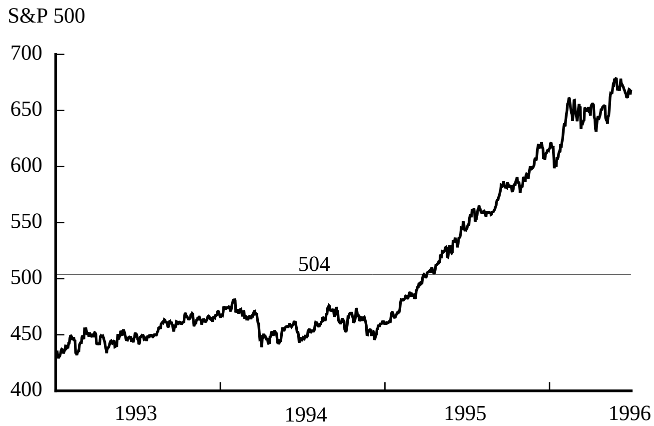

Notice that the S&P 500 index does not fluctuate around the sample mean of 504, but rather has a distinct trend upwards. Comparing the S&P 500 series to the simulated nonstationary data in Chart 4.3, we see that it has all the markings of a nonstationary process.

4.3 Investigating the random-walk model

Thus far we have focused on a simple version of the random walk model (Eq. [4.15]) to demonstrate some important time series properties of financial (log) prices. Recall that this model describes how the prices of financial assets evolve over time, assuming that logarithmic price changes are identically and independently distributed (IID). These assumptions imply:

1. At each point in time, t, log price changes are distributed with a mean 0 and variance σ^2 (identically distributed). This implies that the mean and variance of the log price changes are **homoskedastic**, or unchanging over time.
2. Log price changes are statistically independent of each other over time (independently distributed). That is to say, the values of returns sampled at different points are completely unrelated

In this section we investigate the validity of these assumptions by analyzing real-world data. We find evidence that the IID assumptions do not hold.[11]

[11] Recent (nonparametric) tests to determine whether a time series is IID are presented in Campbell and Dufour (1995).

4.3.1 Is the distribution of returns constant over time?

Visual inspection of real-world data can be a useful way to help understand whether the assumptions of IID returns hold. Using a time series of returns, we investigate whether the first assumption of IID, identically distributed returns, is indeed valid. We find that it is violated and present the following data as evidence.

Charts 4.6 and 4.7 show time series plots of continuously compounded returns for the USD/DEM and USD/FRF exchange rates, respectively.[12]

Chart 4.6
USD/DEM returns

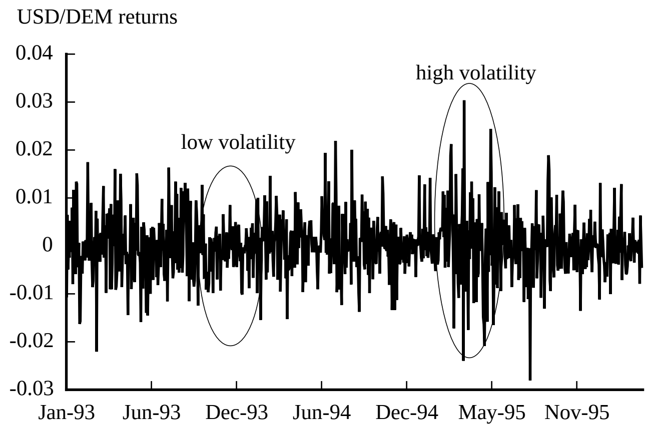

Chart 4.7
USD/FRF returns

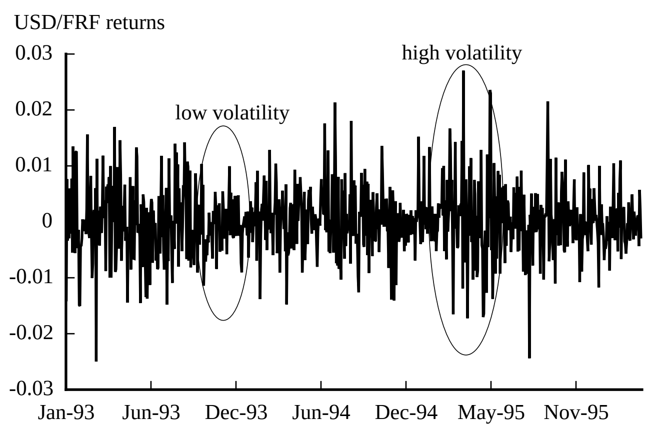

These time series show clear evidence of volatility clustering. That is, periods of large returns are clustered and distinct from periods of small returns, which are also clustered. If we measure such volatility in terms of variance (or its square root, i.e., the standard deviation), then it is fair to think that variance changes with time, reflecting the clusters of large and small returns. In terms of the model in Eq. [4.15], this means that σ_t^2 is changing with time (t). In statistics, changing variances are often denoted by the term **heteroscedasticity**.

[12] This notation (i.e., USD per DEM) is not necessarily market convention.

In Charts 4.6 and 4.7 we also notice not only the individual volatility clustering, but the correlation of the clusters **between** return series. For example, note that periods of high volatility in USD/DEM returns coincide with high volatility in USD/FRF returns. Such correlation between returns series motivates the development of multivariate models, that is, models of returns that measure not only individual series variance (volatility), but also the correlation **between** return series.

4.3.2 Are returns statistically independent over time?

Having established, albeit informally, the possibility of time-varying variances, and consequently, a violation of the identically distributed assumption, we now investigate the validity of the independence assumption, i.e., the second assumption of IID. From our methods and the data that we present in the following sections (4.3.2.1 through 4.3.2.3), we conclude that returns in a given series are **not** independent of each other.

In Charts 4.6 and 4.7, the persistence displayed by the volatility clusters shows some evidence of autocorrelation in variances. That is, the variances of the series are correlated across time. If returns are statistically independent over time, then they are not autocorrelated. Therefore, a natural method for determining if returns are statistically independent is to test whether or not they are autocorrelated. In order to do so, we begin by defining correlation and a method of testing for autocorrelation.

4.3.2.1 Autocorrelation of daily log price changes

For a given time series of returns, the autocorrelation coefficient measures the correlation of returns across time. In general, the standard correlation coefficient between two random variables *X* and *Y* is given by the covariance between *X* and *Y* divided by their standard deviations:

[4.23] $$\rho_{xy} = \frac{\sigma^2_{xy}}{\sigma_x \sigma_y}$$

where σ^2_{xy} represents the covariance between *X* and *Y*. A simple way to understand what covariance measures is to begin with the definition of variance. The variance of a random variable *X* is a measure of the variation of *X* around its mean, μ_X. The mathematical expression for variance is

[4.24] $$E\left[(X - \mu_X)^2\right]$$

where the term *E*[ ] is the mathematical expectation—or more simply, the average. Whereas the variance measures the magnitude of variation of one random variable (in this case *X*), covariance measures the covariation of two random variables (say, *X* and *Y*). It follows that if the variance of *X* is the expected value of $(X - \mu_X)$ times $(X - \mu_X)$, then the covariance of *X* and *Y* is the expected value of $(X - \mu_X)$ times $(Y - \mu_Y)$, or

[4.25] $$E\left[(X - \mu_X)(Y - \mu_Y)\right]$$

Now, for a time series of observations $r_t, t = 1 \ldots T$, the *k*th order autocorrelation coefficient $\rho(k)$ is defined as:

[4.26] $$\rho_k = \frac{\sigma^2_{t,t-k}}{\sigma_t \sigma_{t-k}} = \frac{\sigma^2_{t,t-k}}{\sigma^2_t}$$

Notice that since $\rho(k)$ operates on just one series the subscripts on the covariance and standard deviation refer to the time index on the return series. For a given sample of returns, $r_t, t = 1 \ldots T$, we can estimate Eq. [4.26] using the sample autocorrelation coefficient which is given by:

$$[4.27] \qquad \hat{\rho}_k = \frac{\sum_{t=k+1}^{T} \{ (r_t - \bar{r})(r_{t-k} - \bar{r}) \} / [T - (k-1)]}{\sum_{t=1}^{T} \{ (r_t - \bar{r})^2 \} / [T - 1]}$$

where k = number of lags (days), and $\bar{r} = \frac{1}{T} \sum_{t=1}^{T} r_t$, is the sample mean.

If a time series is not autocorrelated then estimates of $\hat{\rho}_k$ will not be significantly different from 0. In fact, when there is a large amount of historical returns available, we can calculate a 95% confidence band around 0 for each autocorrelation coefficient[13] as $\pm \frac{1.96}{\sqrt{T}}$.

Charts 4.8 and 4.9 show the sample autocorrelation coefficient $\hat{\rho}_k$ plotted against different lags k (measured in days), along with the 95% confidence band around zero for USD/DEM foreign exchange and S&P 500 log price changes, respectively, for the period January 4, 1990 to June 24, 1996. These charts are known as correlograms. The dashed lines represent the upper and lower 95% confidence bands $\pm 4.7\%$. If there is no autocorrelation, that is, if the series are purely random, then we expect only one in twenty of the sample autocorrelation coefficients to lie outside the confidence bands.

Chart 4.8
Sample autocorrelation coefficients for USD/DEM foreign exchange returns

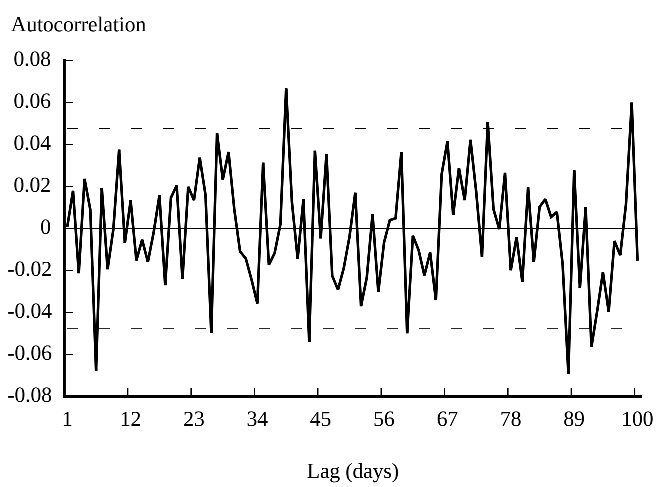

[13] This an asymptotic test statistic since it relies on a large value of T, say, $T > 1000$. See Harvey (p. 43, 1993).

Chart 4.9
Sample autocorrelation coefficients for USD S&P 500 returns

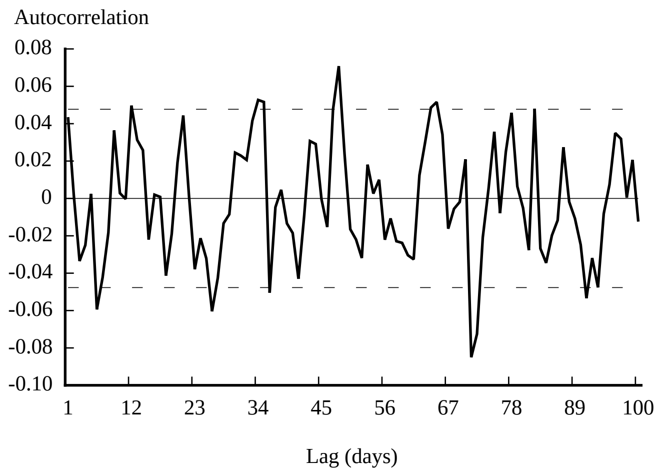

Overall, both charts show very little evidence of autocorrelation in daily log price changes. Even in the cases where the autocorrelations are outside the confidence bands, the autocorrelation coefficients are quite small (less than 10%).

4.3.2.2 Box-Ljung statistic for daily log price changes

While the above charts are useful for getting a general idea about the level of autocorrelation of log price changes, there are more formal methods of testing for autocorrelation. An often cited method is the Box-Ljung (BL) test statistic,[14] defined as

$$[4.28] \qquad BL(p) = T \cdot (T+2) \sum_{k=1}^{p} \frac{\rho_k^2}{T-k}$$

Under the null hypothesis that a time series is not autocorrelated, $BL(p)$, is distributed chi-squared with p degrees of freedom. In Eq. [4.28], p denotes the number of autocorrelations used to estimate the statistic. We applied this test to the USD/DEM and S&P 500 returns for $p = 15$. In this case, the 5% chi-squared critical value is 25. Therefore, values of the $BL(10)$ statistic greater than 25 implies that there is statistical evidence of autocorrelation. The results are shown in Table 4.3.

Table 4.3
Box-Ljung test statistic

Series	$\hat{BL}(15)$
USD/DEM	15
S&P 500	25

[14] See West and Cho (1995) for modifications to this statistic.

We also applied this test to the daily log price changes of a selected series of commodity futures contracts because, when plotted against time, these series appear autocorrelated. In these tests we chose p = 10 which implies a critical value of 18.31 at the 95% confidence level. Table 4.4 presents the results along with the first order autocorrelation coefficient, ρ_1 .

Table 4.4
Box-Ljung statistics

Contract*	**Maturity (mths.)**	$\hat{\rho}_1$	$\hat{BL}(10)$
WTI	1	–0.0338	5.24
	3	–0.0586	7.60
	6	–0.0927	13.62
	12	–0.1323	25.70
LME			
Copper	3	–0.0275	8.48
	15	–0.0900	19.04
	27	–0.1512	16.11

* Note that the higher autocorrelation associated with contracts with longer maturities may be due to the fact that such contracts are less liquid than contracts with short maturities.

The preceding tests show little evidence of autocorrelation for some daily log price change series. The fact that the autocorrelation is not strong agrees with previous research. It is often found that financial returns over the short-run (daily) are autocorrelated but the magnitudes of the autocorrelation are too small (close to zero) to be economically significant.[15] For longer return horizons (i.e., beyond a year), however, there is evidence of significant negative autocorrelation (Fama and French, 1988).

4.3.2.3 Autocorrelation of squared daily log price changes (returns)

As previously stated, although returns (log price changes) are uncorrelated, they may not be independent. In the academic literature, such dependence is demonstrated by the autocorrelation of the variances of returns. **Alternatively expressed, while the returns are not autocorrelated, their squares are autocorrelated**. And since the expected value of the squared returns are variances[16], autocorrelation in the squared returns implies autocorrelation in variances. The relationship between squared returns and variances is evident from the definition of variance, σ_t^2 .

$$\begin{aligned} \sigma_t^2 &= E[r_t - E(r_t)]^2 \\ &= E\left(r_t^2\right) - [E(r_t)]^2 \end{aligned} \qquad [4.29]$$

Assuming that the mean of the returns is zero, i.e., $E(r_t) = 0$, we get $\sigma_t^2 = E\left(r_t^2\right)$.

[15] In other words, it would be very difficult to form profitable trading rules based on autocorrelation in daily log price changes (Tucker, 1992). Also, more recent work has shown that over short horizons, autocorrelation in daily returns may be the result of institutional factors rather than purely inefficient markets (Boudoukh, Richardson and Whitelaw, 1994).

[16] This is true if the expected values of returns are zero.The plausibility of assuming a mean of zero for daily returns will be discussed in Section 5.3.1.1.

Charts 4.10 and 4.11 show time series of squared returns for the USD/DEM exchange rate and for the S&P 500 index.

Chart 4.10
USD/DEM returns squared

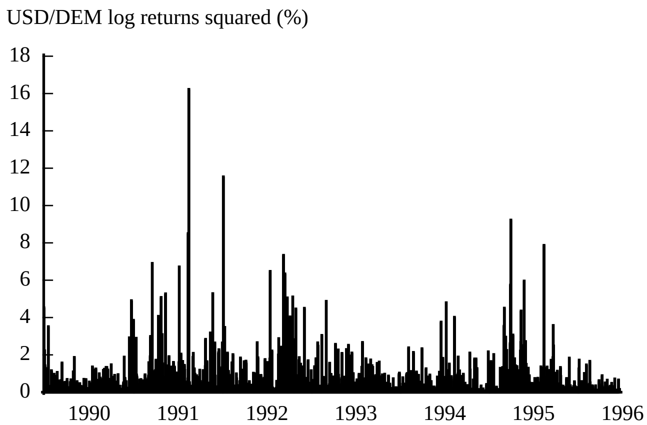

Chart 4.11
S&P 500 returns squared

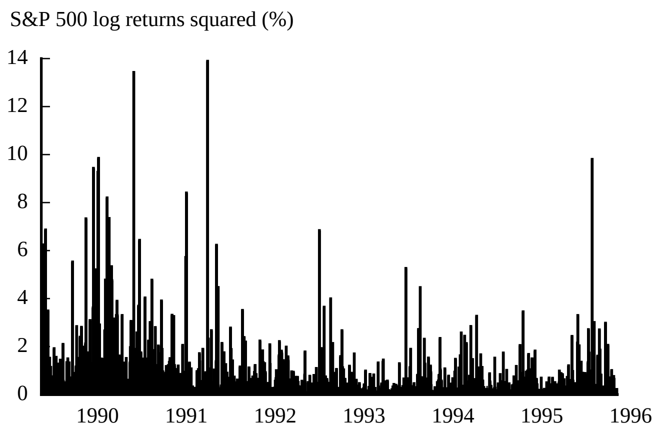

Notice the clusters of large and small spikes in both series. These clusters represent periods of high and low volatility recognized in Section 4.2.1. To analyze the autocorrelation structure of the squared returns, as in the case of log price changes, we compute sample autocorrelation coefficients and the Box-Ljung statistic. Charts 4.12 and 4.13 present correlograms for the squared return series of USD/DEM foreign exchange and S&P 500, respectively.

Chart 4.12
Sample autocorrelation coefficients of USD/DEM squared returns

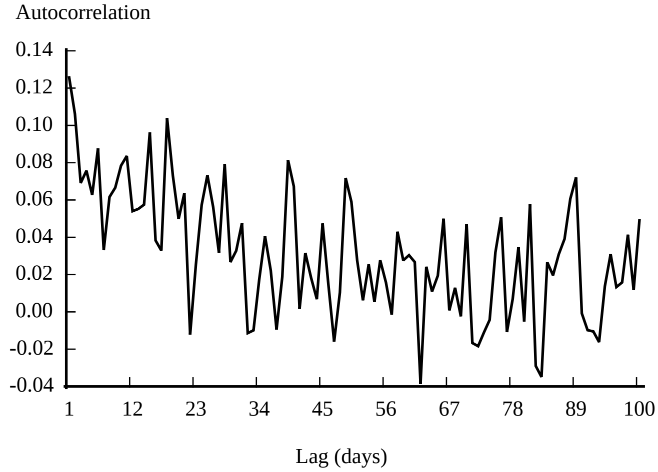

Chart 4.13
Sample autocorrelation coefficients of S&P 500 squared returns

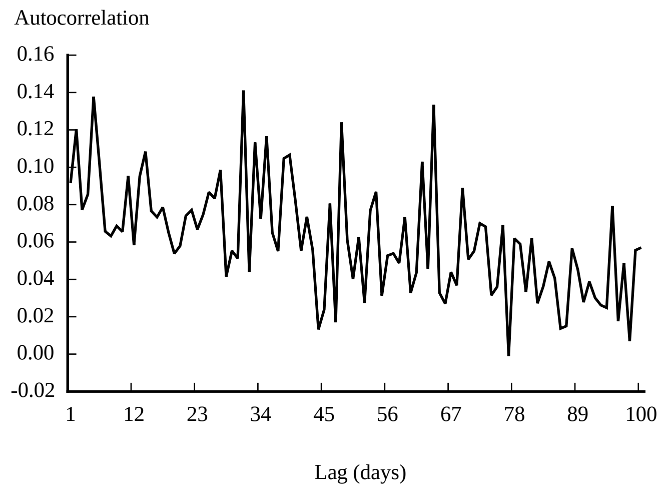

Comparing the correlograms (Charts 4.8 and 4.9) based on daily log price changes to those based on the **squared** daily log price changes (Charts 4.12 and 4.13), we find the autocorrelation coefficients of the squared log price changes are larger and more persistent than those for log price changes. In fact, much of the significant autocorrelation in the squared log price changes is positive and well above the asymptotic 95% confidence band of 4.7%.[17] The Box-Ljung statistics for the squared log price change series are presented in Table 4.5.

Table 4.5
Box-Ljung statistics on squared log price changes (cv = 25)

Series	$\hat{BL}(15)$
USD/DEM	153
S&P 500	207

This table shows the dramatic effect that the squared log price changes has on the BL test. For all three series we reject the null hypothesis that the variances of daily returns are not autocorrelated.[18]

4.3.3 Multivariate extensions

Thus far, we have focused our attention on the empirical properties of individual returns time series. It appears that the variances of returns that were analyzed vary with time and are autocorrelated. As stated in Section 4.3.1, returns appear correlated (through their variances, at least) not only across time but also across securities. The latter finding motivates a study of the empirical properties of correlation, or more precisely, covariance between two return series.

We investigate whether covariances are autocorrelated by using the same logic applied to variances. Recall that we determined whether variances are autocorrelated by checking whether observed squared returns are autocorrelated. We used Eq. [4.29] to show the relation between variances and squared returns. Now, suppose we are interested in the covariance between two return series $r_{1,t}$ and $r_{2,t}$. We can derive a relationship between the covariance, $\sigma^2_{12,t}$, and observed returns as follows. We begin with a definition of covariance between $r_{1,t}$ and $r_{2,t}$.

$$\begin{aligned} \sigma^2_{12,t} &= E\{[r_{1,t} - E(r_{1,t})][r_{2,t} - E(r_{2,t})]\} \\ &= E(r_{1,t}r_{2,t}) - E(r_{1,t})E(r_{2,t}) \end{aligned} \qquad [4.30]$$

Assuming that the mean of the returns is zero for both return series, we get

$$\sigma^2_{12,t} = E(r_{1,t}r_{2,t}) \qquad [4.31]$$

In words, Eq. [4.31] states that the covariance between $r_{1,t}$ and $r_{2,t}$ is the expectation of the cross-product of returns minus the product of the expectations. In models explaining variances, the focus is often on squared returns because of the presumption that for daily returns, squared expected returns are small. Focusing on cross-products of returns can be justified in the same way.

[17] Note that this confidence band may not be appropriate due to the fact that the underlying data are not returns, but squared returns.

[18] For a discussion on tests of autocorrelation on squared returns (residuals) see McLeod and Li (1983) and Li and Mak (1994).

Chart 4.14 presents a time series of the cross product ($r_{1,t}$ times $r_{2,t}$) of the returns on USD/DEM and USD/FRF exchange rates. This series is a proxy for the covariance between the returns on the two exchange rates.

Chart 4.14
Cross product of USD/DEM and USD/FRF returns

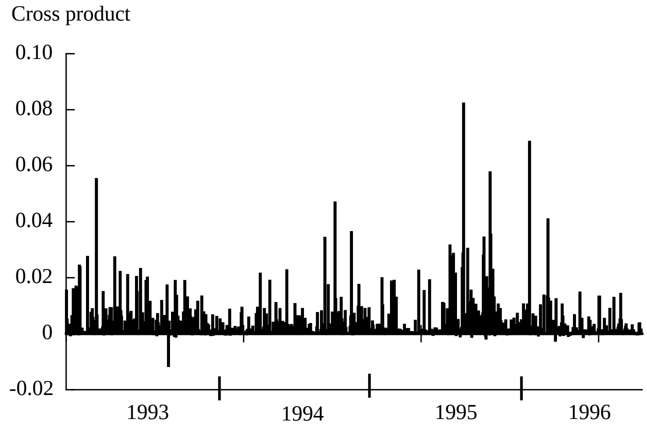

Chart 4.14 shows that the covariance (correlation) between the returns on the two exchange rates is positive over a large segment of the sample period. Time series generated from the cross product of two return series not only offers insight into the temporal dynamics of correlation but also can be used in a regression context to determine the stability of correlations over time.

Similar to the correlogram of squared returns, the correlogram of the cross product of returns on the two exchange rates can be used to determine whether the covariance of these two series are autocorrelated. Chart 4.15 shows the autocorrelations of the cross-products of returns on USD/DEM and USD/FRF exchange rates plotted against 50 daily lags.

Chart 4.15
Correlogram of the cross product of USD/DEM and USD/FRF returns

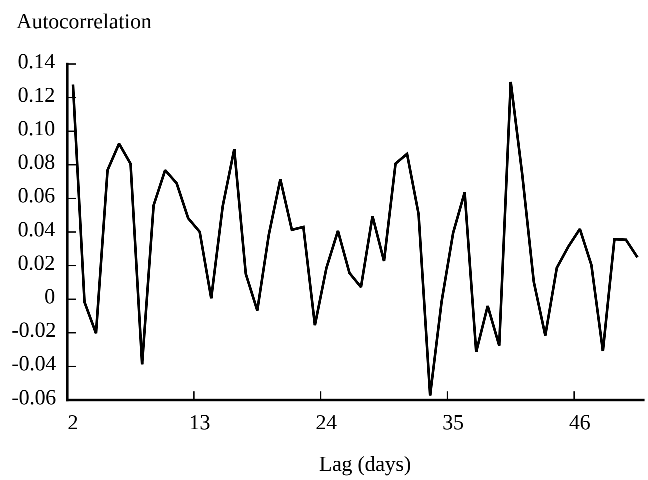

The *BL*(10) test associated with the cross product of returns on the two exchange rate series is 37, which is statistically significant (i.e., there is evidence of autocorrelation) at the 95% confidence level.

4.4 Summary of our findings

Up to this point, Chapter 4 focused on the dynamic features of daily continuously compounded returns, otherwise known as log price changes, and developed the topic as follows:

- We introduced three versions of the random walk model to describe how financial prices evolve over time. We used a particular version of this model (Eq. [4.15]) to highlight the differences between stationary (mean-reverting) and nonstationary time series.

- We investigated the assumptions that log price changes are identically and independently distributed.

  - To determine whether the distribution that generates returns is identical over time, we plotted log price changes against time. From time series plots of returns and their squares we observed the well documented phenomenon of "volatility clustering" which implies that the variance of daily log price changes vary over time (i.e., they are heteroscedastic), thus violating the identical assumption.[19]

  - To test independence, we analyzed the autocorrelation coefficients of both log price changes and squared log price changes. We found that while daily log price changes have small autocorrelations, their squares often have significant autocorrelations.

Much of this analysis has focused on short-horizon (daily) returns. In general, however, observed distributions of returns with longer horizons, such as a month or a quarter, are often different from distributions of daily returns.[20]

From this point, Chapter 4 reviews how returns are assumed to be distributed at each point in time. Specifically, we describe the normal distribution in detail. In RiskMetrics, it is assumed that returns are distributed according to the conditional normal distribution.

4.5 A review of historical observations of return distributions

As shown in Eq. [4.15] and Eq. [4.17], returns were assumed to follow, respectively, an unconditional and conditional normal distribution. The implications of the assumption that financial returns are normally distributed, at least unconditionally, has a long history in finance. Since the early work of Mandelbrot (1963) and Fama (1965), researchers have documented certain stylized facts about the statistical properties of financial returns. A large percentage of these studies focus on high frequency or daily log price changes. Their conclusions can be summarized in four basic observations:

- Financial return distributions have "fat tails." This means that extreme price movements occur more frequently than implied by a normal distribution.

- The peak of the return distribution is higher and narrower than that predicted by the normal distribution. Note that this characteristic (often referred to as the "thin waist") along with fat tails is a characteristic of a **leptokurtotic** distribution.

---

[19] See for example, Engle and Bollerslev (1986).

[20] See, for example, Richardson and Smith (1993)

- Returns have small autocorrelations.
- Squared returns often have significant autocorrelations.

Chart 4.16 illustrates a leptokurtotic distribution of log price changes in USD/DEM exchange rates for the period 28-Mar-96 through 12-Apr-96 and compares it to a normal distribution. In this chart, the leptokurtotic distribution can be thought of as a smoothed histogram, since it is obtained through a smoothing process known as "kernel density estimation."[21] A kernel density estimate of the histogram, rather than the histogram itself, is often used since it produces a smooth line that is easier to compare to the true density function (normal, in this example).

Chart 4.16
Leptokurtotic vs. normal distribution

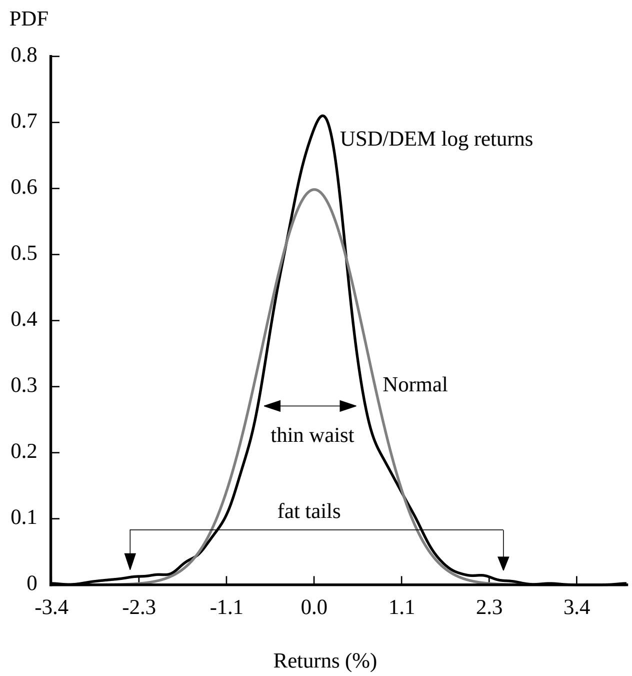

4.5.1 Modeling methods

Having documented the failure of the normal distribution to accurately model returns, researchers started looking for alternative modeling methods, which have since evolved into two classes: **unconditional** (time-independent) and **conditional distributions** (time-dependent) of returns.

Models in the class of unconditional distribution of returns assume that returns are independent of each other and that the return-generating process is linear with parameters that are independent of past realizations. An example of a model that falls into this class is the standard normal distribution with mean μ and variance σ^2 (note there is **no** time subscript). Other examples of unconditional distribution models include infinite-variance symmetric and asymmetric stable Paretian distributions, and finite variance distributions including the t-distribution, mixed-diffusion-jump model, and the compound normal model.

[21] See Silverman (1986).

The second class of models, the conditional distribution of returns, arises from evidence that refutes the identically and independently distributed assumptions (as presented in Sections 4.3.1 and 4.3.2). Models in this category, such as the GARCH and Stochastic Volatility, treat volatility as a time-dependent, persistent process. These models are important because they account for volatility clustering, a frequently observed phenomenon among return series.

The models for characterizing returns are presented in Table 4.6 along with supporting references.

Table 4.6
Model classes

Distribution	Model		Reference
Unconditional (time independent)	Infinite variance:	symmetric stable Paretian	Mandelbrot (1963)
		asymmetric stable Paretian	Tucker (1992)
	Finite variance:	Normal	Bachelier (1900)
		Student t	Blattberg & Gonedes (1974)
		Mixed diffusion jump	Jorion (1988)
		Compound normal	Kon (1988)
Conditional (time dependent)	GARCH:	Normal	Bollerslev (1986)
		Student t	Bollerslev (1987)
	Stochastic Volatility:	Normal	Ruiz (1994)
		Student t	Harvey et. al (1994)
		Generalized error distribution	Ruiz (1994)

It is important to remember that while conditional and unconditional processes are based on different assumptions, except for the unconditional normal model, models from both classes generate data that possess fat tails.[22]

4.5.2 Properties of the normal distribution

All of the models presented in Table 4.6 are parametric in that the underlying distributions depend on various parameters. One of the most widely applied parametric probability distribution is the normal distribution, represented by its "bell shaped" curve.

This section reviews the properties of the normal distribution as they apply to the RiskMetrics method of calculating VaR. Recall that the VaR of a single asset (at time *t*) can be written as follows:

$$VaR_t = [1 - \exp(-1.65\sigma_{t|t-1})]\, V_{t-1} \quad [4.32]$$

or, using the common approximation

$$VaR_t \cong 1.65\sigma_{t|t-1} V_{t-1} \quad [4.33]$$

where V_{t-1} is the marked-to-market value of the instrument and $\sigma_{t|t-1}$ is the standard deviation of continuously compounded returns for time *t* made at time *t*–1.

[22] For a specific comparison between time-dependent and time-independent processes, see Ghose and Kroner (1993).

4.5.2.1 Mean and variance

If it is assumed that returns are generated according to the normal distribution, then it is believed that the entire distribution of returns can be characterized by two parameters: its mean and variance. Mathematically, the normal probability density function for a random variable r_t is[23]

$$f(r_t) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp -\left(\frac{1}{2\sigma^2}\right)(r_t - \mu)^2 \qquad [4.34]$$

where

μ = mean of the random variable, which affects the location of the distribution′s peak

σ^2 = variance of the random variable, which affects the distribution′s width

$\pi \cong 3.1416$

Note that the normal distribution as shown in Eq. [4.34] is an unconditional distribution since the mean and variance parameters are not time-dependent and, therefore, do not have time subscripts.

Chart 4.17 shows how the mean and variance affect the shape of the normal distribution.

Chart 4.17
Normal distribution with different means and variances

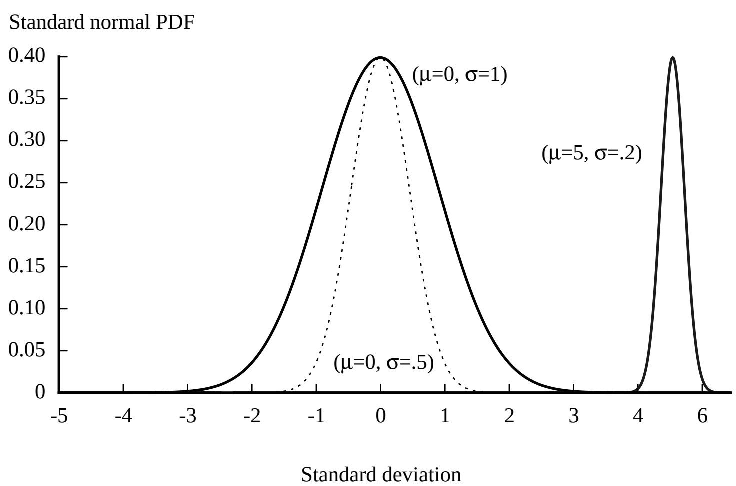

Now that we have an understanding of the role of the mean and variance in the normal distribution we can present their formulae. The mathematical expression for the mean and variance of some random variable r_t, are as follows:

$$\mu = E[r_t] \quad \text{(mean)} \qquad [4.35]$$
$$\sigma^2 = E\left[(r_t - \mu)^2\right] \quad \text{(variance)}$$

[23] Note that we are abusing notation since r_t represents both a random variable and observed return. We hope that by the context in which r_t is used it will be clear what we are referring to.

where E[ ] denotes the mathematical expectation. Two additional measures that we will make reference to within this document are known as skewness and kurtosis. Skewness characterizes the asymmetry of a distribution around its mean. The expression for skewness is given by

[4.36] $$s^3 = E\left[(r_t - \mu)^3\right] \text{ (skewness)}$$

For the normal distribution skewness is zero. In practice, it is more convenient to work with the skewness coefficient which is defined as

[4.37] $$\gamma = \frac{E\left[(r_t - \mu)^3\right]}{\sigma^3} \text{ (skewness coefficient)}$$

Kurtosis measures the relative peakedness or flatness of a given distribution. The expression for kurtosis is given by

[4.38] $$s^4 = E\left[(r_t - \mu)^4\right] \text{ (kurtosis)}$$

As in the case of skewness, in practice, researchers frequently work with the kurtosis coefficient defined as

[4.39] $$\kappa = \frac{E\left[(r_t - \mu)^4\right]}{\sigma^4} \text{ (kurtosis coefficent)}$$

For the normal distribution, kurtosis is 3. This fact leads to the definition of excess kurtosis which is defined as kurtosis minus 3.

4.5.2.2 Using percentiles to measure market risk

Market risk is often measured in terms of a **percentile** (also referred to as **quantile**) of a portfolio’s return distribution. The attractiveness of working with a percentile rather than say, the variance of a distribution, is that a percentile corresponds to both a magnitude (e.g., the dollar amount at risk) and an exact probability (e.g., the probability that the magnitude will not be exceeded).

The *p*th percentile of a distribution of returns is defined as the value that exceeds *p* percent of the returns. Mathematically, the pth percentile (denoted by α) of a continuous probability distribution, is given by the following formula

[4.40] $$p = \int_{-\infty}^{\alpha} f(r)\, dr$$

where $f(r)$ represents the PDF (e.g., Eq. [4.34])

So for example, the 5th percentile is the value (point on the distribution curve) such that 95 percent of the observations lie above it (see Chart 4.18).

When we speak of percentiles they are often of the percentiles of a **standardized distribution,** which is simply a distribution of mean-centered variables scaled by their standard deviation. For example, suppose the log price change r_t is normally distributed with mean μ_t and variance σ_t^2. The standardized return $\tilde{r}_t$ is defined as

[4.41] $$\tilde{r}_t = \frac{r_t - \mu_t}{\sigma_t}$$

Therefore, the distribution of $\tilde{r}_t$ is normal with mean 0 and variance 1. An example of a standardized distribution is presented above ($\mu = 0$, $\sigma = 1$). Chart 4.18 illustrates the positions of some selected percentiles of the standard normal distribution.[24]

Chart 4.18
Selected percentile of standard normal distribution

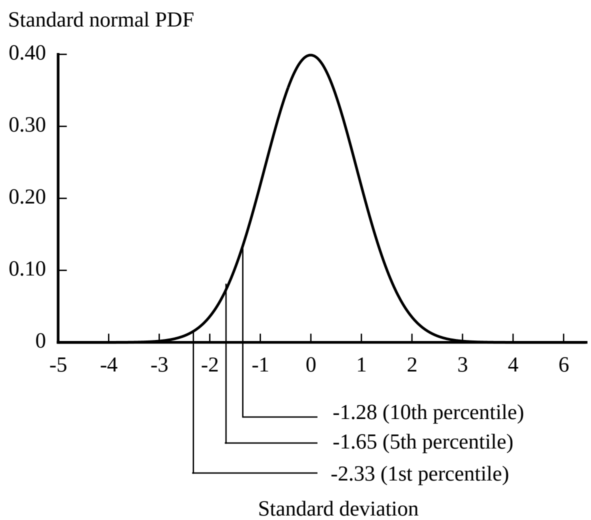

We can use the percentiles of the standard distribution along with Eq. [4.41] to derive the percentiles of observed returns. For example, suppose that we want to find the 5th percentile of r_t, under the assumption that returns are normally distributed. We know, by definition, that

[4.42a] $$\text{Probability}\,(\tilde{r}_t < -1.65) = 5\%$$

[4.42b] $$\text{Probability}\,[\,(r_t - \mu_t)/\sigma_t < -1.65] = 5\%$$

From Eq. [4.42b], re-arranging terms yields

[4.43] $$\text{Probability}\,(r_t < -1.65\sigma_t + \mu_t) = 5\%$$

According to Eq. [4.43], there is a 5% probability that an observed return at time *t* is less than –1.65 times its standard deviation plus its mean. Notice that when $\mu_t = 0$, we are left with the standard result that is the basis for short-term horizon VaR calculation, i.e.,

[4.44] $$\text{Probability}\,(r_t < -1.65\sigma_t) = 5\%$$

[24] Note that the selected percentiles above (1%, 5%, and 10%) reside in the *tails* of the distribution. Roughly, the tails of a distribution are the areas where less then, say, 10% of the observations fall.

4.5.2.3 One-tailed and two-tailed confidence intervals

Equation [4.44] is very important as the basis of VaR calculations in RiskMetrics. It should be recognized, however, that there are different ways of stating the confidence interval associated with the same risk tolerance. For example, since the normal distribution is **symmetric**, then

[4.45]
$$\text{Probability}\,(r_t < -1.65\sigma_t + \mu_t) = \text{Probability}\,(r_t > 1.65\sigma_t + \mu_t) = 5\%$$

Therefore, since the entire area under the probability curve in Chart 4.18 is 100%, it follows that

[4.46a]
$$\text{Probability}\,(-1.65\sigma_t + \mu_t < r_t < 1.65\sigma_t + \mu_t) = 90\%$$

[4.46b]
$$\text{Probability}\,(-1.65\sigma_t + \mu_t < r_t) = 95\%$$

Charts 4.19 and 4.20 show the relationship between a one-tailed 95% confidence interval and a two-tailed 90% confidence interval. Notice that the statements in Eqs. [4.46a] and [4.46b] are consistent with Eq. [4.45], a 5% probability that the return being less than –1.65 standard deviations.[25]

Chart 4.19
One-tailed confidence interval

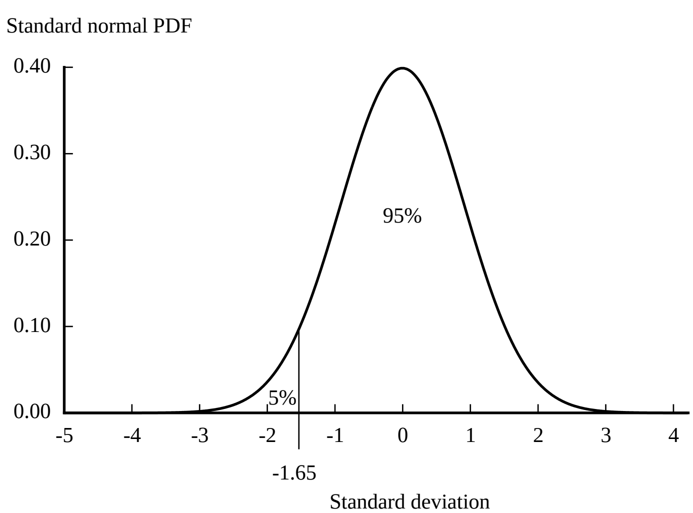

[25] The two statements are not equivalent in the context of formal hypothesis testing. See DeGroot (1989, chapter 8).

Chart 4.20
Two-tailed confidence interval

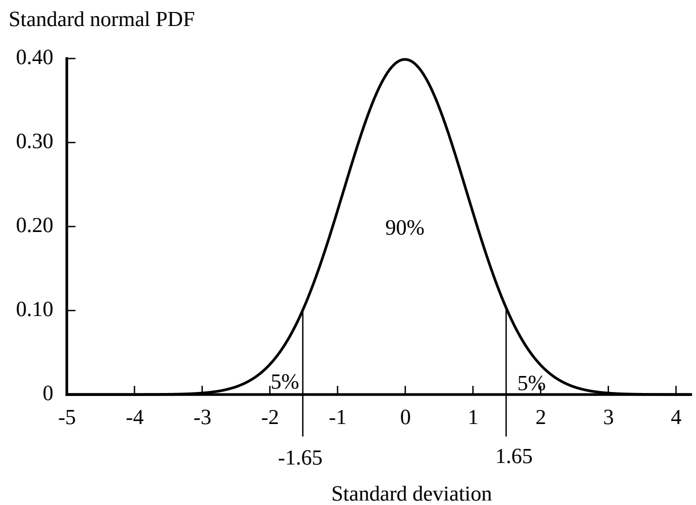

Table 4.7 shows the confidence intervals that are prescribed by standard and BIS-compliant versions of RiskMetrics, and at which the one-tailed and two-tailed tests yield the same VaR figures.[26]

Table 4.7
VaR statistics based on RiskMetrics and BIS/Basel requirements

	Confidence interval	
RiskMetrics method	**One-tailed**	**Two-tailed**
Standard	95% (-1.65σ)	90% ($-/+1.65\sigma$)
BIS/Basel Regulatory	99% (-2.33σ)	98% ($-/+2.33\sigma$)

4.5.2.4 Aggregation in the normal model
An important property of the normal distribution is that the sum of normal random variables is itself normally distributed.[27] This property is useful since portfolio returns are the weighted sum of individual security returns.

As previously stated (p. 49) RiskMetrics assumes that the return on a portfolio, $r_{p,t}$, is the weighted sum of N underlying returns (see Eq. [4.12]). For practical purposes we require a model of returns that not only relates the underlying returns to one another but also relates the distribution of the weighted sum of the underlying returns to the portfolio return distribution. To take an example, consider the case when $N = 3$, that is, the portfolio return depends on three underlying returns. The portfolio return is given by

[4.47] $$r_{pt} = w_1 r_{1,t} + w_2 r_{2,t} + w_3 r_{3,t}$$

[26] For ease of exposition we ignore time subscripts.

[27] These random variables must be drawn from a multivariate distribution.

We can model each underlying return as a random walk that is similar to Eq. [4.17]. This yields

[4.48a] $$r_{1,t} = \mu_1 + \sigma_{1,t}\varepsilon_{1,t}$$

[4.48b] $$r_{2,t} = \mu_2 + \sigma_{2,t}\varepsilon_{2,t}$$

[4.48c] $$r_{3,t} = \mu_2 + \sigma_{3,t}\varepsilon_{3,t}$$

Now, since we have three variables we must account for their movements relative to one another. These movements are captured by pairwise correlations. That is, we define measures that quantify the linear association between each pair of returns. Assuming that the ε_t's are multivariate normally (MVN) distributed we have the model

[4.49] $$\begin{bmatrix}\varepsilon_{1,t}\\ \varepsilon_{2,t}\\ \varepsilon_{3,t}\end{bmatrix} \sim MVN\left(\begin{bmatrix}0\\0\\0\end{bmatrix}, \begin{bmatrix}1 & \rho_{12,t} & \rho_{13,t}\\ \rho_{21,t} & 1 & \rho_{23,t}\\ \rho_{31,t} & \rho_{32,t} & 1\end{bmatrix}\right), \text{ or more succinctly, } \varepsilon \sim MVN(\mu_t, R_t)$$

where parameter matrix R_t represents the correlation matrix of $(\varepsilon_{1,t}, \varepsilon_{2,t}, \varepsilon_{3,t})$. Therefore, if we apply the assumptions behind Eq. [4.49] (that the sum of MVN random variables is normal) to the portfolio return Eq. [4.47], we know that r_{pt} is normally distributed with mean $\mu_{p,t}$ and variance $\sigma^2_{p,t}$. The formulae for the mean and variance are

[4.50a] $$\mu_{p,t} = w_1\mu_1 + w_2\mu_2 + w_3\mu_3$$

[4.50b] $$\sigma^2_{p,t} = w_1^2\sigma^2_{p,t} + w_2^2\sigma^2_{p,t} + w_3^2\sigma^2_{p,t} + 2w_1w_2\sigma^2_{12,t} + 2w_1w_3\sigma^2_{13,t} + 2w_2w_3\sigma^2_{23,t}$$

where the terms $\sigma^2_{ij,t}$ represent the covariance between returns i and j. In general, these results hold for $(N \geq 1)$ underlying returns. Since the underlying returns are distributed conditionally multivariate normal, the portfolio return is univariate normal with a mean and variance that are simple functions of the underlying portfolio weights, variances and covariances.

4.5.3 The lognormal distribution

In Section 4.2.1 we claimed that if log price changes are normally distributed, then price, P_t, conditional on P_{t-1} is lognormally distributed. This statement implies that P_t, given P_{t-1}, is drawn from the probability density function

[4.51] $$f(P_t) = \frac{1}{P_{t-1}\sigma_t\sqrt{2\pi}}\exp\left[\frac{-(\ln P_{t-1} - \mu)^2}{2\sigma_t^2}\right] \qquad P_{t-1} > 0$$

where P_t follows a lognormal distribution with a mean and variance given by

[4.52] $$E[P_t] = \exp\left(\mu + 5\sigma_t^2\right)$$

[4.53] $$V(P_t) = \left[\exp 2\mu_t \cdot \exp\left(2\sigma_t^2\right) - \exp\left(\sigma_t^2\right)\right]$$

Chart 4.21 shows the probability density function for the lognormal random variable P_t when $\mu_t = 0, \sigma_t = 1$ and $P_{t-1} = 1$.

Chart 4.21
Lognormal probability density function

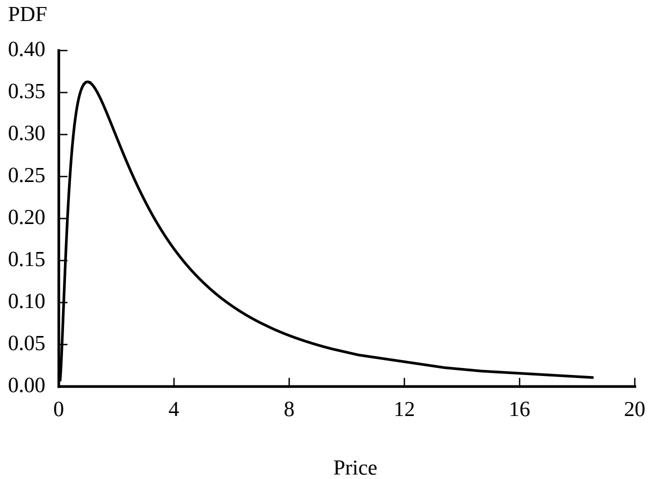

Unlike the normal probability density function, the lognormal PDF has a lower bound greater than zero and is skewed to the right.

4.6 RiskMetrics model of financial returns: A modified random walk

We can now use the results of the last four sections to write down a model of how returns are generated over time. Our analysis has shown that:

- Return variances are heteroscedastic (change over time) and autocorrelated.
- Return covariances are autocorrelated and possess dynamic features.
- The assumption that returns are normally distributed is useful because of the following:

  (i) only the mean and variance are required to describe the entire shape of the distribution[28]

  (ii) the sum of multivariate normal returns is normally distributed. This fact facilitates the description of portfolio returns, which are the weighted sum of underlying returns.

Given these points, we can now state the assumptions underlying the RiskMetrics variance/covariance methodology. Consider a set of N securities, $i = 1\ldots, N$. The RiskMetrics model assumes that returns are generated according to the following model

[4.54]
$$\begin{aligned} r_{i,t} &= \sigma_{i,t}\varepsilon_{i,t} \qquad \varepsilon_{it} \sim N(0,1) \\ \varepsilon_t &\sim MVN(0, R_t) \qquad \varepsilon_t = [\varepsilon_{1t}, \varepsilon_{2t}, \ldots, \varepsilon_{Nt}] \end{aligned}$$

[28] The covariances are also required when there is more than one return series.

where R_t is an *N*x*N* time-dependent correlation matrix. The variance of each return, $\sigma^2_{i,t}$ and the correlation between returns, $\rho_{ij,t}$, are a function of time. The property that the distribution of returns is normal given a time dependent mean and correlation matrix assumes that returns follow a conditional normal distribution—conditional on time. Notice that in Eq. [4.54] we excluded term μ_i. As will be discussed in more detail in Section 5.3.1.1, the mean return represented by μ_i is set to zero.

In Appendix A we propose a set of statistical tests to assess whether observed financial returns follow a conditional normal distribution. In Appendix B we discuss alternative distributions that relax the normality assumption.

4.7 Summary

In this chapter, we presented the statistical and probability assumptions on the evolution and distribution of financial returns in some simple models. This discussion served as background to the specification of the assumptions behind the RiskMetrics VaR methodology.

In review, this chapter covered the following subjects. The chapter began by outlining a simple version of the VaR calculation. We then:

- Defined absolute price change, relative price change, log price change, and returns.
- Showed the importance of understanding the use of different price change definitions.
- Established that RiskMetrics measures changes in portfolio value in terms of continuously-compounded returns.
- Introduced temporal aggregation and cross-section aggregation to show the implications of working with relative and log returns.
- Introduced the random walk model for:[29]
  - Single-price assets
  - Fixed income instruments
- Found evidence that contradicts the assumption that returns are IID (identically and independently) normal. In reality, continuously compounded returns are:
  - Not identical over time. (The variance of the return distribution changes over time)
  - Not statistically independent of each other over time. (Evidence of autocorrelation between return series and within a return series.)
- Explained the properties of the normal distribution, and, lastly,
- Presented the RiskMetrics model as a modified random walk that assumes that returns are conditionally normally distributed.

[29] While the random walk model serves as the basis for many popular models of returns in finance, another class of models that has received considerable attention lately is based on the phenomenon of **long-range dependence**. Briefly, such models are built on the notion that observations recorded in the distant past are correlated to observations in the distant future. (See Campbell, et. al (1995) for a review of long-range dependence models.)

Chapter 5. Estimation and forecast

Chapter 5. Estimation and forecast

Peter Zangari
Morgan Guaranty Trust Company
Risk Management Research
(1-212) 648-8641
zangari_peter@jpmorgan.com

In this chapter we present a methodology for forecasting the parameters of the multivariate conditional normal distribution, i.e., variances and covariances of returns whose empirical properties were examined in Chapter 4, "Statistical and probability foundations." The reason for forecasting variances and covariances of returns is to use them to forecast a portfolio's change in value over a given horizon, which can run over one day to several months.

This chapter is organized as follows:

- Section 5.1 briefly explains why RiskMetrics forecasts of variances and covariances are generated from historical data rather than derived from option prices.
- Section 5.2 describes the RiskMetrics forecasting methodology, i.e.,
  - Use of the exponentially weighted moving average (EWMA) model to produce forecasts of variances and covariances. This includes an explanation as to why the EWMA is preferred to the simple moving average model.
  - How to compute forecasts over longer time horizons, such as one month.

  Section 5.2 also discusses alternative, more advanced methods for forecasting variances and covariances.
- Section 5.3 explains two important implementation issues involving the RiskMetrics forecasts: (1) the reliability of the forecasts in relation to the number of historical data points used to produce them, and (2) the choice of the "decay factor" used in the EWMA model.
- Section 5.4 concludes the chapter with a review of the RiskMetrics forecasting model.

Finally, practitioners often refer to the term "volatility" when speaking of movements in financial prices and rates. In what follows we use the term volatility to mean the standard deviation of continuously compounded financial returns.

5.1 Forecasts from implied versus historical information

RiskMetrics forecasts are based on historical price data, although in theory, they may be derived from option prices.

From a practical point of view, implied forecasts introduce a number of problems. For example, an implied volatility (IV) is based entirely on expectations given a particular option pricing model. Therefore, as noted in Kroner, Kneafsey and Claessens (1995), since most option pricing models assume that the standard deviation is constant, the IV becomes difficult to interpret and will not lead to good forecasts if the option formula used to derive it is not correctly specified. Moreover, IV forecasts are associated with a fixed forecast horizon. For example, the implied volatility derived from a 3 month USD/DEM option is exclusively for a 3 month forecast horizon. However, a risk manager may be interested in the VaR of this option over the next day.

If RiskMetrics were to use implied statistics, it would require observable options prices on all instruments that compose a portfolio. Currently, the universe of consistently observable options prices is not large enough to provide a complete set of implied statistics; generally only exchange-traded options are reliable sources of prices. In particular, the number of implied correlations that can be derived from traded option prices is insignificant compared to the number of correlations required to estimate risks in portfolios consisting of many types of assets.

Academic research has compared the forecasting ability of implied and historical volatility models. The evidence of the superior forecasting ability of historical volatility over implied volatility is mixed, depending on the time series considered. For example, Xu and Taylor (1995, p. 804) note that, "prior research concludes that volatility predictors calculated from options prices are better predictors of future volatility than standard deviations calculated from historical asset price data." Kroner, Kneafsey and Claessens (1995, p. 9), on the other hand, note that researchers are beginning to conclude that GARCH (historical based) forecasts outperform implied volatility forecasts. Since implied standard deviation captures market expectations and pure time series models rely solely on past information, these models can be combined to forecast the standard deviation of returns.

5.2 RiskMetrics forecasting methodology

RiskMetrics uses the exponentially weighted moving average model (EWMA) to forecast variances and covariances (volatilities and correlations) of the multivariate normal distribution. This approach is just as simple, yet an improvement over the traditional volatility forecasting method that relies on moving averages with fixed, equal weights. This latter method is referred to as the simple moving average (SMA) model.

5.2.1 Volatility estimation and forecasting[1]

One way to capture the dynamic features of volatility is to use an exponential moving average of historical observations where the latest observations carry the highest weight in the volatility estimate. This approach has two important advantages over the equally weighted model. First, volatility reacts faster to shocks in the market as recent data carry more weight than data in the distant past. Second, following a shock (a large return), the volatility declines exponentially as the weight of the shock observation falls. In contrast, the use of a simple moving average leads to relatively abrupt changes in the standard deviation once the shock falls out of the measurement sample, which, in most cases, can be several months after it occurs.

For a given set of T returns, Table 5.1 presents the formulae used to compute the equally and exponentially weighted (standard deviation) volatility.

Table 5.1
Volatility estimators*

Equally weighted	**Exponentially weighted**
$\sigma = \sqrt{\frac{1}{T}\sum_{t=1}^{T}(r_t - \bar{r})^2}$	$\sigma = \sqrt{(1-\lambda)\sum_{t=1}^{T}\lambda^{t-1}(r_t - \bar{r})^2}$

* In writing the volatility estimators we intentionally do not use time subscripts.

In comparing the two estimators (equal and exponential), notice that the exponentially weighted moving average model depends on the parameter λ ($0 < \lambda < 1$) which is often referred to as the **decay factor**. This parameter determines the relative weights that are applied to the observations (returns) and the effective amount of data used in estimating volatility. Ways of estimating λ are discussed in detail in Section 5.3.2.

[1] In this section we refer loosely to the terms estimation and forecast. The reader should note, however, that these terms do have distinct meanings.

We point out that in writing the EWMA estimator in Table 5.1 we applied the approximation

$$\sum_{j=1}^{T} \lambda^{j-1} \cong \frac{1}{(1-\lambda)} \tag{5.1}$$

These two expressions are equivalent in the limit, i.e., as $T \to \infty$. Moreover, for purpose of comparison to the equally weighted factor 1/*T*, the more appropriate version of the EWMA is

$$\lambda^{t-1} \Big/ \sum_{j=1}^{T} \lambda^{j-1} \tag{5.2}$$

rather than $(1-\lambda)\lambda^{t-1}$. Also, notice that when $\lambda = 1$, Eq. [5.2] collapses to 1/*T*.

Charts 5.1 and 5.2 highlight an important difference between equally and exponentially weighted volatility forecasts using as an example the GBP/DEM exchange rate in the fall of 1992. In late August of that year, the foreign exchange markets went into a turmoil that led a number of Europe’s currencies to leave the ERM and be devalued. The standard deviation estimate using an exponential moving average rapidly reflected this state of events, but also incorporated the decline in volatility over subsequent months. The simple 6-month moving average estimate of volatility took longer to register the shock to the market and remained higher in spite of the fact that the foreign exchange markets calmed down over the rest of the year.

Chart 5.1
DEM/GBP exchange rate

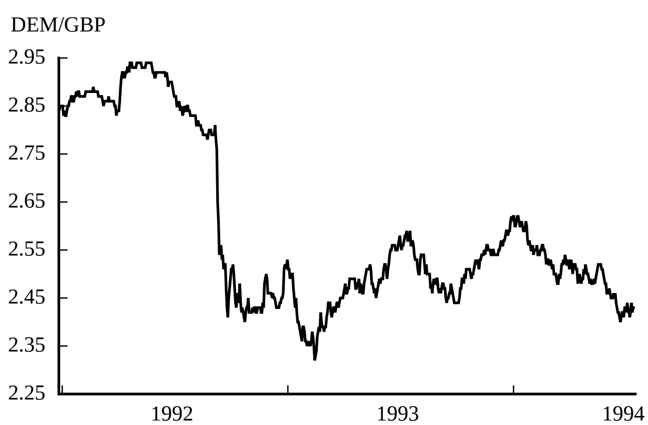

Chart 5.2
Log price changes in GBP/DEM and VaR estimates (1.65σ)

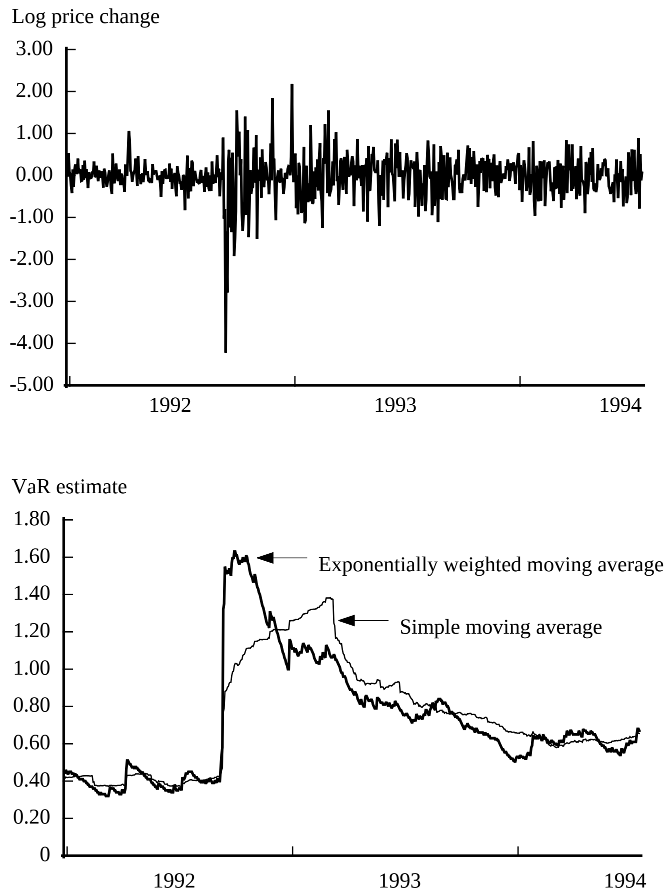

This example would suggest that EWMA is more satisfactory, given that when combined with frequent updates, it incorporates external shocks better than equally weighted moving averages, thus providing a more realistic measure of current volatility.

Although the exponentially weighted moving average estimation ranks a level above simple moving averages in terms of sophistication, it is not complex to implement. To support this point, Table 5.2 presents an example of the computation required to estimate equally and exponentially weighted moving average volatilities. Volatility estimates are based on 20 daily returns on the USD/DEM exchange rate. We arbitrarily choose $\lambda = 0.94$ and keep matters simple by setting the sample mean, $\bar{r}$, to zero.

Table 5.2
Calculating equally and exponentially weighted volatility

	A	B	C	D	Volatility	
Date	**Return USD/DEM (%)**	**Return squared (%)**	**Equal weight (T = 20)**	**Exponential weight (λ = 0.94)**	**Equally weighted, B × C**	**Exponentially weighted, B × D**
28-Mar-96	0.634	0.402	0.05	0.019	0.020	0.007
29-Mar-96	0.115	0.013	0.05	0.020	0.001	0.000
1-Apr-96	-0.460	0.211	0.05	0.021	0.011	0.004
2-Apr-96	0.094	0.009	0.05	0.022	0.000	0.000
3-Apr-96	0.176	0.031	0.05	0.024	0.002	0.001
4-Apr-96	-0.088	0.008	0.05	0.025	0.000	0.000
5-Apr-96	-0.142	0.020	0.05	0.027	0.001	0.001
8-Apr-96	0.324	0.105	0.05	0.029	0.005	0.003
9-Apr-96	-0.943	0.889	0.05	0.030	0.044	0.027
10-Apr-96	-0.528	0.279	0.05	0.032	0.014	0.009
11-Apr-96	-0.107	0.011	0.05	0.034	0.001	0.000
12-Apr-96	-0.160	0.026	0.05	0.037	0.001	0.001
15-Apr-96	-0.445	0.198	0.05	0.039	0.010	0.008
16-Apr-96	0.053	0.003	0.05	0.041	0.000	0.000
17-Apr-96	0.152	0.023	0.05	0.044	0.001	0.001
18-Apr-96	-0.318	0.101	0.05	0.047	0.005	0.005
19-Apr-96	0.424	0.180	0.05	0.050	0.009	0.009
22-Apr-96	-0.708	0.501	0.05	0.053	0.025	0.027
23-Apr-96	-0.105	0.011	0.05	0.056	0.001	0.001
24-Apr-96	-0.257	0.066	0.05	0.060	0.003	0.004
			Standard deviation:	Equally weighted	0.393	
				Exponentially weighted		0.333

Notice that the difference between the two estimated standard deviations results from the different weighting schemes. Whereas the equally weighted approach weights each squared return by 5%, the exponentially weighted scheme applies a 6% weight to the most recent squared return and 1.9% weight to the most distant observation.

An attractive feature of the exponentially weighted estimator is that it can be written in **recursive** form which, in turn, will be used as a basis for making volatility forecasts. In order to derive the recursive form, it is assumed that an infinite amount of data are available. For example, assuming again that the sample mean is zero, we can derive the period $t + 1$ variance forecast, given data available at time t (one day earlier) as

[5.3] $$\sigma^2_{1,\,t+1|t} = \lambda\sigma^2_{1,\,t|t-1} + (1-\lambda)\,r^2_{1,\,t}$$

The 1-day RiskMetrics volatility forecast is given by the expression

[5.4] $$\sigma_{1,\,t+1|t} = \sqrt{\lambda\sigma^2_{1,\,t|t-1} + (1-\lambda)\,r^2_{1,\,t}}$$

The subscript "$t+1|t$" is read "the time $t+1$ forecast given information up to and including time t." The subscript "$t|t-1$" is read in a similar fashion. This notation underscores the fact that we are treating the variance (volatility) as time-dependent. The fact that this period's variance forecast depends on last period's variance is consistent with the observed autocorrelation in squared returns discussed in Section 4.3. We derive Eq. [5.3] as follows.

$$\begin{aligned} \sigma^2_{1,t+1|t} &= (1-\lambda)\sum_{i=0}^{\infty}\lambda^i r^2_{1,t-i} \\ &= (1\text{-}\lambda)\left(r^2_{1,t} + \lambda r^2_{1,t-1} + \lambda^2 r^2_{1,t-2} + \ldots\right) \\ &= (1\text{-}\lambda)\, r^2_{1,t} + \lambda(1-\lambda)\left(r^2_{1,t-1} + \lambda r^2_{1,t-2} + r^2_{1,t-3}\right) \\ &= \lambda\sigma^2_{1,t|t-1} + (1-\lambda)\, r^2_{1,t} \end{aligned} \tag{5.5}$$

Using daily returns, Table 5.3 presents an example of how Eq. [5.3] can be used in practice to produce a 1-day volatility forecast on USD/DEM returns for the period March 28 through April 24, 1996.

Table 5.3
Applying the recursive exponential weighting scheme to compute volatility
Daily returns on USD/DEM

Date	A: Return USD/DEM	B: Recursive variance	Date	A: Return USD/DEM	B: Recursive variance
28-Mar-96	0.633	0.401	11-Apr-96	–0.107	0.296
29-Mar-96	0.115	0.378	12-Apr-96	–0.159	0.280
1-Apr-96	–0.459	0.368	15-Apr-96	–0.445	0.275
2-Apr-96	0.093	0.346	16-Apr-96	0.053	0.258
3-Apr-96	0.176	0.327	17-Apr-96	0.152	0.244
4-Apr-96	–0.087	0.308	18-Apr-96	–0.318	0.236
5-Apr-96	–0.142	0.291	19-Apr-96	0.424	0.232
8-Apr-96	0.324	0.280	22-Apr-96	–0.708	0.248
9-Apr-96	–0.943	0.316	23-Apr-96	–0.105	0.234
10-Apr-96	–0.528	0.314	24-Apr-96	–0.257	0.224

*Initial variance forecast = initial return squared. Figures following this number are obtained by applying the recursive formula.

The volatility forecast made on April 24 for the following day is the square root of 0.224% (the variance) which is 0.473%.

5.2.1.1 Covariance and correlation estimation and forecasts

We use the EWMA model to construct covariance and correlation forecasts in the same manner as we did volatility forecasts except that instead of working with the square of one series, we work with the product of two different series. Table 5.4 presents covariance estimators based on equally and exponentially weighted methods.

Table 5.4
Covariance estimators

Equally weighted	**Exponentially weighted**
$\sigma_{12}^2 = \frac{1}{T}\sum_{t=1}^{T}(r_{1t}-\bar{r}_1)(r_{1t}-\bar{r}_2)$	$\sigma_{12}^2 = (1-\lambda)\sum_{j=1}^{T}\lambda^{j-1}(r_{1t}-\bar{r}_1)(r_{1t}-\bar{r}_2)$

Analogous to the expression for a variance forecast (Eq. [5.3]), the covariance forecast can also be written in recursive form. For example, the 1-day covariance forecast between any two return series, $r_{1,t}$ and $r_{2,t}$ made at time t is

[5.6] $$\sigma^2_{12,\,t+1|t} = \lambda\sigma^2_{12,\,t|t-1} + (1-\lambda)\,r_{1t}\cdot r_{2t}$$

We can derive Eq. [5.6] as follows.

[5.7] $$\begin{aligned}\sigma^2_{1,\,t+1|t} &= (1-\lambda)\sum_{i=0}^{\infty}\lambda^i r_{1,\,t-i}\cdot r_{2,\,t-i}\\ &= (1-\lambda)\left(r_{1,\,t}\cdot r_{2,\,t} + \lambda r_{1,\,t-1}\cdot r_{2,\,t-1} + \lambda^2 r_{2,\,t-2}\cdot r_{2,\,t-2} + \ldots\right)\\ &= (1-\lambda)\,r_{1,\,t}\cdot r_{2,\,t}\\ &\quad + \lambda(1-\lambda)\left(r_{1,\,t-1}\cdot r_{2,\,t-1} + \lambda r_{1,\,t-2}\cdot r_{2,\,t-2} + \lambda^2 r_{1,\,t-3}\cdot r_{2,\,t-3}\right)\\ &= \lambda\sigma^2_{12,\,t|t-1} + (1-\lambda)\,r_{1,\,t-1}\cdot r_{2,\,t-1}\end{aligned}$$

In order to derive correlation forecasts we apply the corresponding covariance and volatility forecast. Recall that correlation is the covariance between the two return series, say, $r_{1,t}$ and $r_{2,t}$, divided by the product of their standard deviations. Mathematically, **the one-day RiskMetrics prediction of correlation is given by the expression**

[5.8] $$\rho_{12,\,t+1|t} = \frac{\sigma^2_{12,\,t+1|t}}{\sigma_{1,\,t+1|t}\sigma_{2,\,t+1|t}}$$

Table 5.5 presents an example of how to compute recursive covariance and correlation forecasts applied to the USD/DEM exchange rate and S&P 500 return series.

Table 5.5
Recursive covariance and correlation predictor

Date	Returns USD/DEM (%)	Returns S&P 500 (%)	Recursive variance USD/DEM	Recursive variance S&P 500	Recursive covariance (λ = 0.94)	Recursive correlation (λ = 0.94)
28-Mar-96	0.634	0.005	0.402	0.000	0.003	1.000
29-Mar-96	0.115	–0.532	0.379	0.017	–0.001	–0.011
1-Apr-96	-0.460	1.267	0.369	0.112	–0.036	–0.176
2-Apr-96	0.094	0.234	0.347	0.109	–0.032	–0.166
3-Apr-96	0.176	0.095	0.328	0.103	–0.029	–0.160
4-Apr-96	-0.088	–0.003	0.309	0.097	–0.028	–0.160
5-Apr-96	-0.142	–0.144	0.291	0.092	–0.025	–0.151
8-Apr-96	0.324	–1.643	0.280	0.249	–0.055	–0.209
9-Apr-96	-0.943	–0.319	0.317	0.240	–0.034	–0.123
10-Apr-96	-0.528	–1.362	0.315	0.337	0.011	0.035
11-Apr-96	-0.107	–0.367	0.296	0.325	0.013	0.042
12-Apr-96	-0.160	0.872	0.280	0.351	0.004	0.012
15-Apr-96	-0.445	0.904	0.275	0.379	–0.020	–0.063
16-Apr-96	0.053	0.390	0.259	0.365	–0.018	–0.059
17-Apr-96	0.152	–0.527	0.245	0.360	–0.022	–0.073
18-Apr-96	-0.318	0.311	0.236	0.344	–0.026	–0.093
19-Apr-96	0.424	0.227	0.233	0.327	–0.019	–0.069
22-Apr-96	-0.708	0.436	0.249	0.318	–0.036	–0.129
23-Apr-96	-0.105	0.568	0.235	0.319	–0.038	–0.138
24-Apr-96	-0.257	–0.217	0.224	0.302	–0.032	–0.124

Note that the starting points for recursion for the covariance is 0.634×0.005. From Table 5.5 we can see that the correlation prediction for the period 24-Apr-96 through 25-Apr-96 is –12.4%.

5.2.2 Multiple day forecasts

Thus far, we have presented 1-day forecasts which were defined over the period t through $t + 1$, where each t represents one business day. Risk managers, however, are often interested in forecast horizons greater than one-day. We now demonstrate how to construct variance (standard deviation) and covariance (correlation) forecasts using the EWMA model over longer time horizons. Generally speaking, the T-period (i.e., over T days) forecasts of the variance and covariance are, respectively,

$$\sigma^2_{1,\, t+T|t} = T\sigma^2_{1,\, t+1|t} \quad \text{or} \quad \sigma_{1,\, t+T|t} = \sqrt{T}\sigma_{1,\, t+1|t} \qquad [5.9]$$

and

$$\sigma^2_{12,\, t+T|t} = T\sigma^2_{12,\, t+1|t} \qquad [5.10]$$

Equations [5.9] and [5.10] imply that the correlation forecasts remain unchanged regardless of the forecast horizon. That is, [5.11] $$\rho_{t+T|t} = \frac{T\sigma^2_{12,t+1|t}}{\sqrt{T}\sigma_{1,t+1|t}\sqrt{T}\sigma_{2,t+1|t}} = \rho_{t+1|t}$$

Notice that multiple day forecasts are simple multiples of one-day forecasts. For example, if we define one month to be equivalent to 25 days, then the 1-month variance and covariance forecasts are 25 times the respective 1-day forecasts and the 1-month correlation is the same as the one-day correlation.[2] We now show how we arrive at Eq. [5.9] and Eq. [5.10].

Recall that RiskMetrics assumes that log prices p_t are generated according to the model

[5.12] $$p_{1,t} = p_{1,t-1} + \sigma_{1,t}\varepsilon_{1,t} \qquad \varepsilon_{1,t} \sim IID\ N(0,1)$$

Recursively solving Eq. [5.12] and writing the model in terms of returns, we get

[5.13] $$r_{1,t+T} = \sum_{s=1}^{T} \sigma_{1,t+s}\varepsilon_{1,t+s}$$

Taking the variance of Eq. [5.13] as of time *t* implies the following expression for the forecast variance

[5.14] $$\sigma^2_{1,t+T} = E_t\left[r^2_{1,t+T}\right] = \sum_{s=1}^{T} E_t\left[\sigma^2_{1,t+s}\right]$$

Similar steps can be used to find the *T* days-ahead covariance forecast, i.e.,

[5.15] $$\sigma^2_{12,t+T} = E_t\left[r_{1,t+T} \cdot r_{2,t+T}\right] = \sum_{s=1}^{T} E_t\left[\sigma^2_{12,t+s}\right]$$

Now, we need to evaluate the right-hand side of Eq. [5.14] and Eq. [5.15]. To do so, we work with the recursive form of the EWMA model for the variance and covariance. To make matters concrete, consider the case where we have two (correlated) return series, $r_{1,t}$ and $r_{2,t}$. In vector form[3], let's write the 1-day forecast of the two variances and covariance as follows:

[2] In RiskMetrics, 1-day and 1-month forecasts differ because we use different decay factors when making the forecasts.

[3] We use the "vec representation" as presented in Engle and Kroner (1995).

[5.16]

$$\sigma^2_{t+1|t} = \begin{bmatrix} \sigma^2_{1,\,t+1|t} \\ \sigma^2_{12,\,t+1|t} \\ \sigma^2_{2,\,t+1|t} \end{bmatrix}$$

$$= \begin{bmatrix} \lambda & 0 & 0 \\ 0 & \lambda & 0 \\ 0 & 0 & \lambda \end{bmatrix} \begin{bmatrix} \sigma^2_{1,\,t|t-1} \\ \sigma^2_{12,\,t|t-1} \\ \sigma^2_{2,\,t|t-1} \end{bmatrix} + \begin{bmatrix} 1-\lambda & 0 & 0 \\ 0 & 1-\lambda & 0 \\ 0 & 0 & 1-\lambda \end{bmatrix} \begin{bmatrix} r^2_{1,\,t} \\ r_{1,\,t} \cdot r_{2,\,t} \\ r^2_{2,\,t} \end{bmatrix}$$

Using the expectation operator at time t, write the forecast over S days as

[5.17]

$$E_t\left[\sigma^2_{t+s}\right] = \begin{bmatrix} \lambda & 0 & 0 \\ 0 & \lambda & 0 \\ 0 & 0 & \lambda \end{bmatrix} \begin{bmatrix} E_t\left[\sigma^2_{1,\,t+s-1}\right] \\ E_t\left[\sigma^2_{12,\,t+s-1}\right] \\ E_t\left[\sigma^2_{2,\,t+s-1}\right] \end{bmatrix}$$

$$+ \begin{bmatrix} 1-\lambda & 0 & 0 \\ 0 & 1-\lambda & 0 \\ 0 & 0 & 1-\lambda \end{bmatrix} \begin{bmatrix} E_t\left[r^2_{1,\,t+s-1}\right] \\ E_t\left[r_{1,\,t+s-1} \cdot r_{2,\,t+s-1}\right] \\ E_t\left[r^2_{2,\,t+s-1}\right] \end{bmatrix}$$

Evaluating the expectations of the squared returns and their cross product yields

[5.18]

$$E_t\left[\sigma^2_{t+s}\right] = \left(\begin{bmatrix} \lambda & 0 & 0 \\ 0 & \lambda & 0 \\ 0 & 0 & \lambda \end{bmatrix} + \begin{bmatrix} 1-\lambda & 0 & 0 \\ 0 & 1-\lambda & 0 \\ 0 & 0 & 1-\lambda \end{bmatrix} \right) \begin{bmatrix} E_t\left[\sigma^2_{1,\,t+s-1}\right] \\ E_t\left[\sigma^2_{12,\,t+s-1}\right] \\ E_t\left[\sigma^2_{2,\,t+s-1}\right] \end{bmatrix}$$

$$= E_t\left[\sigma^2_{t+s-1}\right]$$

That is, the variance forecasts for two consecutive periods are the same. Consequently, the T-period forecast is defined as

[5.19]

$$\sigma^2_{t+T|t} = \sum_{s=1}^{T} E_t\left[\sigma^2_{t+s}\right]$$

$$= T \cdot E_t\left[\sigma^2_{t+1}\right]$$

so that the T-period forecast of the variance/covariance vector is

[5.20]

$$\sigma^2_{t+T|t} = T \cdot \sigma^2_{t+1|t}$$

This leads to the "square root of time" relationship for the standard deviation forecast

[5.21] $$\sigma_{1,\,t+T|t} = \sqrt{T} \cdot \sigma_{1,\,t+1|t}$$

Having found that volatility and covariance forecasts scale with time, a few points are worth noting about Eq. [5.21]. Typically, the “square root of time rule” results from the assumption that variances are constant. Obviously, in the above derivation, volatilities and covariances vary with time. Implicitly, what we are assuming in modeling the variances and covariances as exponentially weighted moving averages is that the variance process is nonstationary. Such a model has been studied extensively in the academic literature (Nelson 1990, Lumsdaine, 1995) and is referred to as the IGARCH model.[4]

In practice, scaling up volatility forecasts may sometimes lead to results that do not make much sense. Three instances when scaling up volatility estimates prove problematic are:

- When rates/prices are mean-reverting (see Section 4.2.3)
- When boundaries limit the potential movements in rates and prices
- When estimates of volatilities optimized to forecast changes over a particular horizon are used for another horizon (jumping from daily to annual forecasts, for example).

Take the simple example of the Dutch guilder to Deutsche mark exchange rate. On March 22, 1995, the cross rate as quoted at London close of business was 1.12048 NLG/DEM. The RiskMetrics daily volatility estimate was 0.1648%, which meant that over the next 24 hours, the rate was likely to move within a 1.1186 to 1.1223 range with 90% probability (the next day’s rate was 1.1213 NLG/DEM).

The Netherlands and Germany have maintained bilateral 2.25% bands within the ERM so scaling up a daily volatility estimate can quickly lead to exchange rate estimates which are extremely unlikely to occur in reality. An example of this is shown by Chart 5.3:

Chart 5.3
NLG/DEM exchange rate and volatility

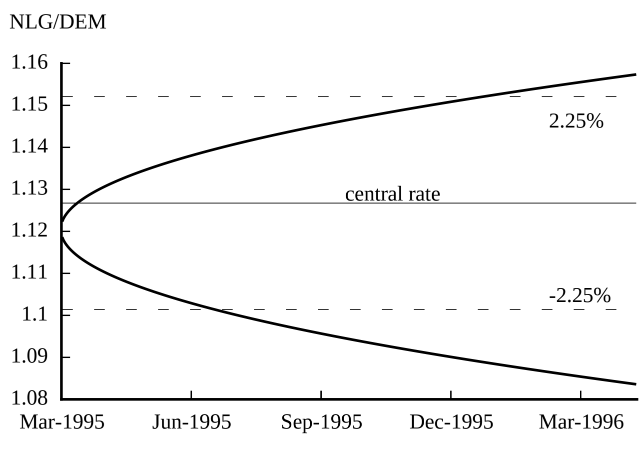

[4] Note that whereas we essentially arrive at a model that reflects an IGARCH (without an intercept), our motivation behind its derivation was more “bottom up” in the sense that we wanted to derive a model that is generally consistent with observed returns while being simple to implement in practice. The formal approach to IGARCH is more “top down” in that a formal statistical model is written down which then maximum likelihood estimation is used to estimate its parameters.

Applying the square root of time rule with caution does not apply exclusively to exchange rates that are constrained by political arrangements. Suppose you had been trying to forecast the S&P 500's potential annual volatility on April 5, 1994. The index stood at 448.3 and its previous declines had increased the daily volatility estimate to 1.39%. Chart 5.4 extends this daily volatility estimate out to the end of the first quarter of 1995 using the square root of time rule. The chart shows how a short term increase in daily volatility would bias an estimate of volatility over any other time horizon, for example, a year.

Chart 5.4
S&P 500 returns and VaR estimates (1.65σ)

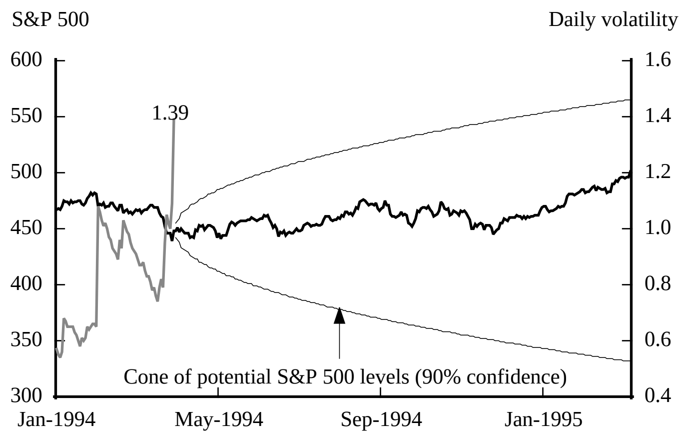

The preceding two examples underscore the importance of understanding how volatility estimates for horizons longer than a day are calculated. When daily volatility forecasts are scaled, nonsensical results may occur because the scale factor does not account for real-world restrictions.

5.2.3 More recent techniques

Research in finance and econometrics has devoted significant efforts in recent years to come up with more formal methods to estimate standard deviations and correlations. These are often referred to as volatility models. The methods range from extreme value techniques (Parkinson, 1980) and two step regression analysis (Davidian and Carroll, 1987), to more complicated nonlinear modelling such as GARCH (Bollerslev, 1986), stochastic volatility (Harvey et. al, 1994) and applications of chaotic dynamics (LeBaron, 1994). Among academics, and increasingly among practitioners, GARCH-type models have gained the most attention. This is due to the evidence that time series realizations of returns often exhibit time-dependent volatility. This idea was first formalized in Engle's (1982) ARCH (Auto Regressive Conditional Heteroscedasticity) model which is based on the specification of conditional densities at successive periods of time with a time-dependent volatility process.

Of the methods just mentioned, the least computationally demanding procedures for estimating volatility are the extreme value and regression methods. Extreme value estimators use various types of data such as high, low, opening and closing prices and transaction volume. While this approach is known for its relative efficiency (i.e., small variance), it is subject to bias. On the other hand, the two step regression method treats the underlying volatility model as a regression involving the absolute value of returns on lagged values. Applications of this method to monthly volatility can be found in Schwert (1989) and Pagan and Schwert (1990).

Since the introduction of the basic ARCH model, extensions include generalized ARCH (GARCH), Integrated GARCH (IGARCH), Exponential GARCH (EGARCH) and Switching

Regime ARCH (SWARCH), just to name a few. Numerous tests of GARCH-type models to foreign exchange and stock markets have demonstrated that these relatively sophisticated approaches can provide somewhat better estimates of volatility than simple moving averages, particularly over short time horizons such as a day or a week.

More recent research on modeling volatility involves Stochastic Volatility (SV) models. In this approach, volatility may be treated as an unobserved variable, the logarithm of which is modeled as a linear stochastic process, such as an autoregression. Since these models are quite new, their empirical properties have yet to be established. However, from a practical point of view, an appealing feature of the SV models is that their estimation is less daunting than their counterpart EGARCH models.[5]

In a recent study, West and Cho (1995) found that GARCH models did not significantly outperform the equally weighted standard deviation estimates in out-of-sample forecasts, except for very short time horizons. In another study on foreign exchange rates and equity returns, Heynen and Kat (1993) showed that while GARCH models have better predictive ability for foreign exchange, the advantage over a simple random walk estimator disappears when the outlook period chosen is more than 20 days.

We have elected to calculate the volatilities and correlations in the RiskMetrics data set using exponential moving averages. This choice is viewed as an optimal balance given the constraints under which most risk management practitioners work.

Since the GARCH models are becoming more popular among practitioners, we demonstrate the behavior of the daily volatility estimator by comparing its forecasts to those produced by a GARCH(1,1) volatility model with normal disturbances. If r_t represents the time t daily return, then the return generating process for the GARCH(1,1) volatility model is given by

$$\begin{aligned} r_t &= \sigma_t \varepsilon_t \qquad \varepsilon_t \sim IID\ N(0,1) \\ \sigma_t^2 &= 0.0147 + 0.881\sigma_{t-1}^2 + 0.0828 r_{t-1}^2 \end{aligned} \tag{5.22}$$

This model is parameterized according to the results produced in Ruiz (1993). They were estimated from daily return data for the British pound. The following graph shows variance forecasts produced by this model and the exponential estimator with the decay factor set to 0.94. The forecasts from the EWMA are based on the following equation:

$$\sigma_{t+1|t}^2 = 0.94\sigma_{t|t-1}^2 + 0.06 r_t^2 \tag{5.23}$$

[5] Bayesian SV models, on the other hand, are computationally intensive.

Chart 5.5
GARCH(1,1)-normal and EWMA estimators
GBP parameters

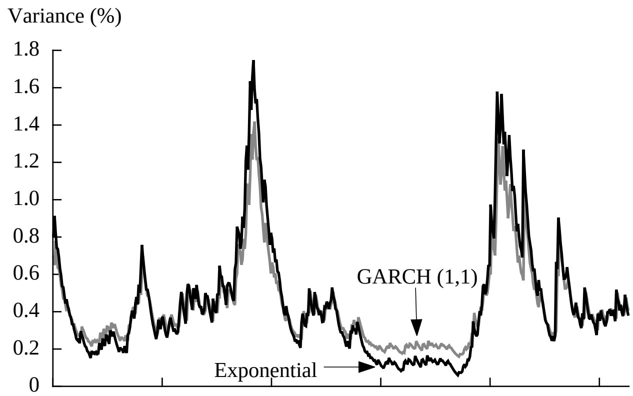

Notice from Chart 5.5, the dynamics of the exponential model's forecasts closely mimic those produced by the GARCH(1,1) model. This should not be surprising given our findings that the exponential model is similar in form to the IGARCH model.

A natural extension of univariate GARCH and Stochastic Volatility models has been to model conditional covariances and correlations. With the ability to estimate more parameters of the return generating process comes growing computational complexity.[6] Often, to make models tractable, restrictions are placed on either the process describing the conditional covariance matrix or the factors that explain covariance dynamics. Recent discussion and applications of multivariate GARCH models include Engle and Kroner (1995), Karolyi (1995), King, Sentena and Wadhwani (1994). Harvey (1993) presents work on multivariate extensions to the stochastic volatility models.

5.3 Estimating the parameters of the RiskMetrics model

In this section we address two important issues that arise when we estimate RiskMetrics volatilities and correlations. The first issue concerns the estimation of the sample mean. In practice, when we make volatility and correlation forecasts we set the sample mean to zero. The second issue involves the estimation of the exponential decay factor which is used in volatility and correlation forecasts.

5.3.1 Sample size and estimation issues

Whenever we must estimate and/or forecast means, standard deviations and correlations, we would like to be reasonably confident in the results. Here, confidence is measured by the standard error of the estimate or forecast; in general, the smaller the standard error, the more confident we are about its value. It is important, therefore, to use the largest samples available when computing these statistics. We illustrate the relationship between sample size and confidence intervals next. For ease of exposition we use equally weighted statistics. The results presented below carry over to the case of exponentially weighted statistics as well.

[6] With respect to the required computation of the bivariate EGARCH model, Braun, Nelson and Sunier (1991) note that, “ease of computation is, alas, not a feature even of the bivariate model. For, example, the FORTRAN code for computing the analytic derivatives … ran to forty pages.”

5.3.1.1 The sample mean

Table 5.6 shows that the mean estimates for USD/DEM foreign exchange returns and S&P 500 returns are –0.114 and –0.010 percent, respectively. To show the variability of the sample mean, Chart 5.6 presents historical estimates of the sample mean for USD/DEM exchange rate returns. Each estimate of the mean is based on a 74-day rolling window, that is, every day in the sample period we estimate a mean based on returns over the last 74 days.

Table 5.6

Mean, standard deviation and correlation calculations

USD/DEM and S&P500 returns

	Returns	
Date	**USD/DEM**	**S&P 500**
28-Mar-96	0.634	0.005
29-Mar-96	0.115	–0.532
1-Apr-96	–0.460	1.267
2-Apr-96	0.094	0.234
3-Apr-96	0.176	0.095
4-Apr-96	–0.088	–0.003
5-Apr-96	–0.142	–0.144
8-Apr-96	0.324	–1.643
9-Apr-96	–0.943	–0.319
10-Apr-96	–0.528	–1.362
11-Apr-96	–0.107	–0.367
12-Apr-96	–0.160	0.872
15-Apr-96	–0.445	0.904
16-Apr-96	0.053	0.390
17-Apr-96	0.152	–0.527
18-Apr-96	–0.318	0.311
19-Apr-96	0.424	0.227
22-Apr-96	–0.708	0.436
23-Apr-96	–0.105	0.568
24-Apr-96	–0.257	–0.217
Mean	–0.114	0.010
Standard deviation	0.393	0.688
Correlation	–0.180	

Chart 5.6
USD/DEM foreign exchange

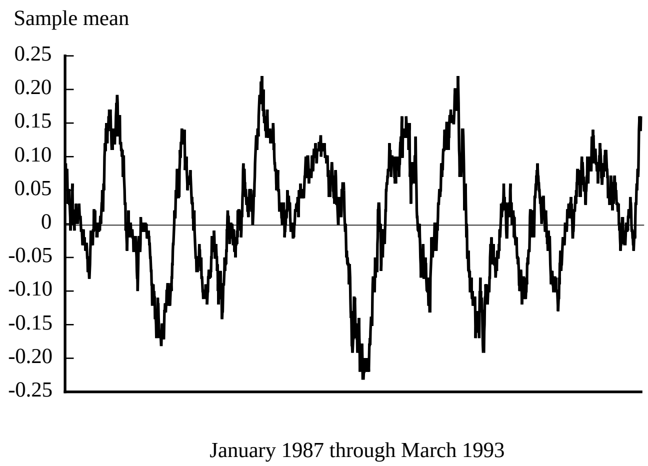

Chart 5.6 shows how the estimates of the mean of returns on USD/DEM fluctuate around zero. An interesting feature of the equally weighted sample mean estimator is that the mean estimate does not depend directly on the number of observations used to construct it. For example, recall that the 1-day log return is defined as $r_t = \ln(P_t / P_{t-1}) = p_t - p_{t-1}$. Now, the sample mean of returns for the period $t = 1,\ldots, T$ is

[5.24]
$$\bar{r} = \frac{1}{T}\sum_{t=1}^{T} p_t - p_{t-1} = \frac{1}{T}(p_T - p_0)$$

Hence, we see that the sample mean estimator depends only on the first and last observed prices; all other prices drop out of the calculation. Since this estimator does not depend on the number of observed prices between $t = 0$ and $t = T$ but rather on the length of the sample period, neither does its standard error. The implication of this effect can best be demonstrated with a simple example.[7]

Suppose a price return has a standard deviation of 10 percent and we have 4 years' of historical price data. The standard deviation of the sample mean is $10/\sqrt{4} = 5$ percent. So, if the average annual return were 20 percent over the 4-year sample (which consists of over 1000 data points), a 95 percent confidence region for the true mean would range from 10 percent to 30 percent.

In addition, recall that the variance of a returns series, r_t, can be written as $\sigma^2 = E(r_t^2) - [E(r_t)]^2$. Jorion (1995) notes that with daily data the "average term $E(r_t^2)$ dominates the term $[E(r_t)]^2$ by a typical factor of 700 to one. Therefore, ignoring expected returns is unlikely to cause a perceptible bias in the volatility estimate."

To reduce the uncertainty and imprecision of the estimated mean, it may be more accurate to set the mean to some value which is consistent with financial theory. **In RiskMetrics, we assume that the mean value of daily returns is zero. That is, standard deviation estimates are cen-**

[7] This example is adapted from Figlewski, (1994).

tered around zero, rather than the sample mean. Similarly, when computing the covariance, deviations of returns are taken around zero rather than the sample mean.

5.3.1.2 Volatility and correlation

Volatility and correlation forecasts based on the EWMA model requires that we choose an appropriate value of the decay factor λ. As a practical matter, it is important to determine the effective number of historical observations that are used in the volatility and correlation forecasts.

We can compute the number of effective days used by the variance (volatility) and covariance (correlation) forecasts. To do so, we use the metric

[5.25] $$\Omega_K^{\infty} = (1-\lambda)\sum_{t=K}^{\infty}\lambda^t$$

Setting Ω_K^{∞} equal to a value —the tolerance level (Υ_L)— we can solve for *K*, the effective number of days of data used by the EWMA. The formula for determining *K* is

[5.26] $$K = \frac{\ln\Upsilon_L}{\ln\lambda}$$

Equation [5.26] is derived as follows

[5.27] $$\Omega_K^{\infty} = (1-\lambda)\sum_{t=K}^{\infty}\lambda^t = \Upsilon_L$$

which implies

[5.28] $$\lambda^K(1-\lambda)\left(1+\lambda+\lambda^2+\ldots\right) = \Upsilon_L$$

Solving Eq. [5.28] for *K* we get Eq. [5.26].

Table 5.7 shows the relationship between the tolerance level, the decay factor, and the effective amount of data required by the EWMA.

Table 5.7
The number of historical observations used by the EWMA model
daily returns

	Days of historical data at tolerance level:			
Decay factor	**0.001%**	**0.01%**	**0.1%**	**1 %**
0.85	71	57	43	28
0.86	76	61	46	31
0.87	83	66	50	33
0.88	90	72	54	36
0.89	99	79	59	40
0.9	109	87	66	44
0.91	122	98	73	49
0.92	138	110	83	55
0.93	159	127	95	63
0.94	186	149	112	74
0.95	224	180	135	90
0.96	282	226	169	113
0.97	378	302	227	151
0.98	570	456	342	228
0.99	1146	916	687	458

For example, setting a tolerance level to 1% and the decay factor to 0.97, we see the EWMA uses approximately 151 days of historical data to forecast future volatility/correlation. Chart 5.7 depicts the relationship between the tolerance level and the amount of historical data implied by the decay factor

Chart 5.7
Tolerance level and decay factor

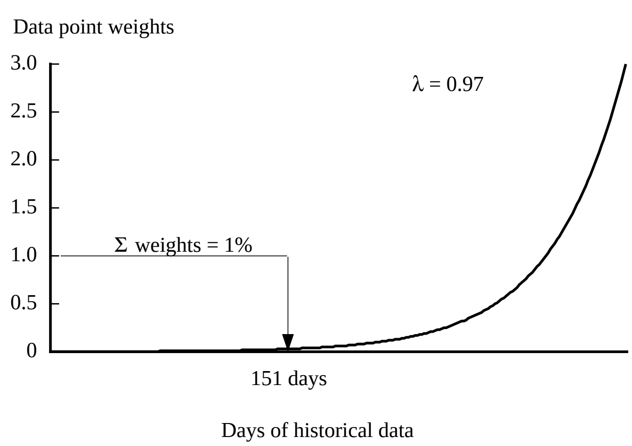

Chart 5.8 shows the relationship between the number of days of data required by EWMA and various values of the decay factor.

Chart 5.8
Relationship between historical observations and decay factor

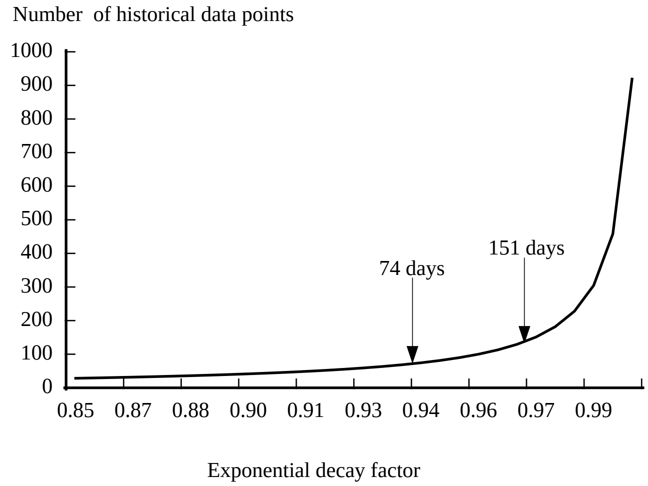

For a different perspective on the relationship between the number of data points used and different values of the decay factor, consider Chart 5.9. It shows the weights for different decay factors over a fixed window size of T = 100 (approximately 6 months' of data).

Chart 5.9
Exponential weights for *T* = 100
decay factors = 1, .99, .97, .95, .93

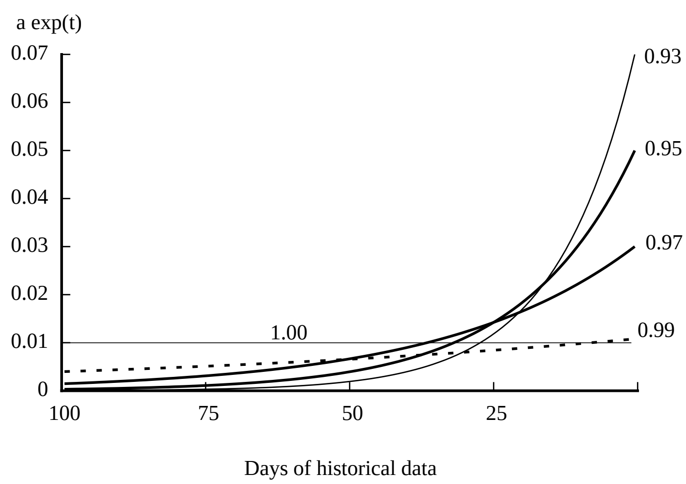

Note that while the decay factor of 0.93 weighs the most recent data more than the factor 0.99, after 40 days, the weight associated with the decay factor of 0.93 is below the weight of 0.99. Hence, the closer the decay factor is to 1, the less responsive it is to the most recent data.

Now we consider the effect of sample size on volatility and correlation forecasts. Chart 5.10 presents two historical time series of 1-day volatility forecasts on the returns series in USD/DEM exchange rate. One volatility series was constructed with a decay factor of 0.85, the other used 0.98. (Refer to Table 5.7 for the relationship between the decay factor and the amount of data used).

Chart 5.10
One-day volatility forecasts on USD/DEM returns
λ = 0.85 (black line), λ = 0.98 (gray line)

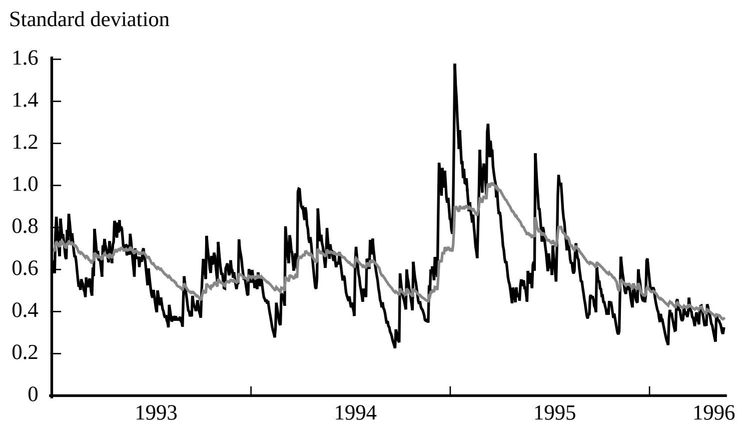

As expected, the volatility forecasts based on more historical observations are smoother than those that rely on much less data.

One-day forecasts of correlation between the returns on the USD/DEM foreign exchange rate and S&P 500 for two different decay factors are presented in Chart 5.11.

Chart 5.11
One-day correlation forecasts for returns on USD/DEM FX rate and on S&P500
λ = 0.85 (black line), λ = 0.98 (gray line)

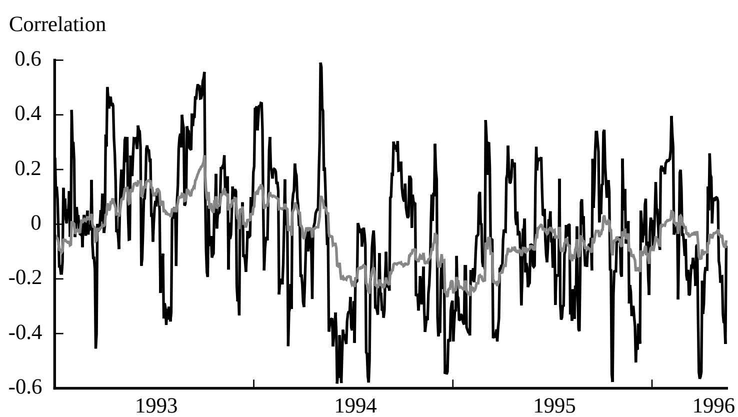

Again, the time series with the higher decay factor produces more stable (though not necessarily more accurate) forecasts.

5.3.2 Choosing the decay factor

In this section we explain how we determine the decay factors (λ's) that are used to produce the RiskMetrics volatility and correlation forecasts. We begin by describing the general problem of choosing 'optimal' λ's for volatilities and correlations that are consistent with their respective covariance matrix. We then discuss **how** RiskMetrics chooses its two optimal decay factors; one for the daily data set (λ = 0.94), and the other for the monthly data set (λ = 0.97).

RiskMetrics produces volatility and correlation forecasts on over 480 time series. This requires 480 variance forecasts and 114,960 covariance forecasts. Since these parameters comprise a covariance matrix, the optimal decay factors for each variance and covariance forecast are not independent of one another. We explain this concept with a simple example that consists of two return series, $r_{1,t}$ and $r_{2,t}$. The covariance matrix associated with these returns is given by

[5.29] $$\Sigma = \begin{bmatrix} \sigma_1^2(\lambda_1) & \sigma_{12}^2(\lambda_3) \\ \sigma_{21}^2(\lambda_3) & \sigma_2^2(\lambda_2) \end{bmatrix}$$

We write each parameter explicitly as a function of its decay factor. As we can see from Eq. [5.29], the covariance matrix, Σ, is a function of 3 decay factors, λ_1, λ_2 and λ_3. Now, Σ, to be properly defined must contain certain properties. For example, Σ must be such that the following three conditions are met:

- The variances, σ_1^2 and σ_2^2, cannot be negative
- The covariances σ_{12}^2 and σ_{21}^2 must be equal (i.e., Σ is symmetric)
- The correlation between $r_{1,t}$ and $r_{2,t}$ has the range $-1 \le \rho \le 1$. (Recall the definition of correlation, ρ, $\rho = \sigma_{12}^2 / (\sigma_1 \sigma_2)$.

It follows then that decay factors must be chosen such that they not only produce good forecasts of future variances and covariances, but that the values of these decay factors are consistent with the properties of the covariance matrix to which they belong.

In theory, while it is possible to choose optimal decays factors that are consistent with their respective covariance matrix, in practice this task is exceedingly complex for large covariance matrices (such as the kind that RiskMetrics produces that has 140,000 elements). Therefore, it becomes necessary to put some structure (restrictions) on the optimal λ's.

RiskMetrics applies one optimal decay factor to the entire covariance matrix. That is, we use one decay factor for the daily volatility and correlation matrix and one for the monthly volatility and correlation matrix. This decay factor is determined from individual variance forecasts across 450 time series (this process will be discussed in Section 5.3.2.2).

Recently, Crnkovic and Drachman (1995)[8] have shown that while it is possible to construct a covariance matrix with different decay factors that is positive semi-definite, this matrix is subject to substantial bias.[9]

We now describe a measure applied by RiskMetrics to determine the optimal decay factor, i.e., that decay factor that provides superior forecast accuracy.

[8] From personal communication.

[9] See Section 8.3 for an explanation of positive semi-definite and its relationship to covariance matrices.

5.3.2.1 Root mean squared error (RMSE) criterion[10]

The definition of the time $t+1$ forecast of the variance of the return, r_{t+1}, made one period earlier is simply $E_t\left[r_{t+1}^2\right] = \sigma_{t+1|t}^2$, i.e., the expected value of the squared return one-period earlier. Similarly, the definition of the time $t+1$ forecast of the covariance between two return series, $r_{1,t+1}$ and $r_{2,t+1}$ made one period earlier is $E_t\left[r_{1,t+1}r_{2,t+1}\right] = \sigma_{12,t+1|t}^2$. In general, these results hold for any forecast made at time $t+j,\ j \geq 1$.

Now, if we define the variance forecast error as $\varepsilon_{t+1|t} = r_{t+1}^2 - \sigma_{t+1|t}^2$ it then follows that the expected value of the forecast error is zero, i.e., $E_t\left[\varepsilon_{t+1|t}\right] = E_t\left[r_{t+1}^2\right] - \sigma_{t+1|t}^2 = 0$. Based on this relation a natural requirement for choosing λ is to minimize average squared errors. When applied to daily forecasts of variance, this leads to the (daily) root mean squared prediction error which is given by

$$[5.30] \qquad RMSE_v = \sqrt{\frac{1}{T}\sum_{t=1}^{T}\left(r_{t+1}^2 - \hat{\sigma}_{t+1|t}^2(\lambda)\right)^2} \quad \text{(variance)}$$

where the forecast value of the variance is written explicitly as a function of λ.

In practice we find the optimal decay factor λ^* by searching for the smallest RMSE over different values of λ. That is, we search for the decay factor that produces the best forecasts (i.e., minimizes the forecast measures).

Although RiskMetrics does not assess the accuracy of covariance forecasts, similar results to those for the variance can be derived for covariance forecasts, i.e., the covariance forecast error is $\varepsilon_{12,t+1|t} = r_{1,t+1}r_{2,t+1} - \sigma_{12,t+1|t}^2$ such that $E_t\left[\varepsilon_{12,t+1|t}\right] = E_t\left[r_{1,t+1}r_{2,t+1}\right] - \sigma_{12,t+1|t}^2 = 0$ and

$$[5.31] \qquad RMSE_c = \sqrt{\frac{1}{T}\sum_{t=1}^{T}\left(r_{1,t+1}r_{2,t+1} - \hat{\sigma}_{12,t+1|t}^2(\lambda)\right)^2} \quad \text{(covariance)}$$

The measures presented above are purely statistical in nature. For risk management purposes, this may not be optimal since other factors come into play that determine the best forecast. For example, the decay factor should allow enough stability in the variance and covariance forecasts so that these forecasts are useful for risk managers who do not update their systems on a daily basis.[11]

Next, we explain how we determine the two RiskMetrics optimal decay factors, one for daily and one for monthly forecasts.

[10] See Appendix C for alternative measures to assess forecast accuracy.

[11] West, Edison and Cho (1993) suggested that an interesting alternative basis for comparing forecasts is to calculate the utility of an investor with a particular utility function investing on the basis of different variance forecasts. We plan to pursue this idea from a risk management perspective in future research.

5.3.2.2 How RiskMetrics chooses its optimal decay factor

RiskMetrics currently processes 480 time series, and associated with each series is an optimal decay factor that minimizes the root mean squared error of the variance forecast (i.e., Eq. [5.30]). We choose RMSE as the forecast error measure criterion.[12] Table 5.8 presents optimal decay factors for return series in five series.

Table 5.8
Optimal decay factors based on volatility forecasts
based on RMSE criterion

Country	Foreign exchange	5-year swaps	10-year zero prices	Equity indices	1-year money market rates
Austria	0.945	—	—	—	—
Australia	0.980	0.955	0.975	0.975	0.970
Belgium	0.945	0.935	0.935	0.965	0.850
Canada	0.960	0.965	0.960	—	0.990
Switzerland	0.955	0.835	—	0.970	0.980
Germany	0.955	0.940	0.960	0.980	0.970
Denmark	0.950	0.905	0.920	0.985	0.850
Spain	0.920	0.925	0.935	0.980	0.945
France	0.955	0.945	0.945	0.985	—
Finland	0.995	—	—	—	0.960
Great Britain	0.960	0.950	0.960	0.975	0.990
Hong Kong	0.980	—	—	—	—
Ireland	0.990	—	0.925	—	—
Italy	0.940	0.960	0.935	0.970	0.990
Japan	0.965	0.965	0.950	0.955	0.985
Netherlands	0.960	0.945	0.950	0.975	0.970
Norway	0.975	—	—	—	—
New Zealand	0.975	0.980	—	—	—
Portugal	0.940	—	—	—	0.895
Sweden	0.985	—	0.980	—	0.885
Singapore	0.950	0.935	—	—	—
United States	—	0.970	0.980	0.980	0.965
ECU	—	0.950	—	—	—

For the daily and monthly data sets we compute one optimal decay factor from the 480+ time series. Denote the *i*th optimal decay factor by $\hat{\lambda}_i$ and let N (i = 1, 2,…, N) denote the number of time series in the RiskMetrics database. Also, let τ_i denote the *i*th RMSE associated with $\hat{\lambda}_i$, i.e., τ_i is the minimum RMSE for the *i*th time series. We derive the one optimal decay factor as follows:

1. Find Π, the sum of all N minimal RMSE's, τ_i's:

$$\Pi = \sum_{i=1}^{N} \tau_i . \qquad [5.32]$$

2. Define the relative error measure:

[12] We have chosen this criterion because it penalizes large forecast errors more severely, and provides more useful results than other common accuracy statistics.

[5.33] $$\theta_i = \tau_i \Big/ \left(\sum_{i=1}^{N} \tau_i \right)$$

3. Define the weight ϕ_i:

[5.34] $$\phi_i = \theta_i^{-1} \Big/ \sum_{i=1}^{N} \theta_i^{-1}$$

where $\sum_{i=1}^{N} \phi_i = 1$

4. The optimal decay factor $\tilde{\lambda}$ is defined as

[5.35] $$\tilde{\lambda} = \sum_{i=1}^{N} \phi_i \hat{\lambda}_i$$

That is, the optimal decay factor applied by RiskMetrics is a weighted average of individual optimal decay factors where the weights are a measure of individual forecast accuracy.

Applying this methodology to both daily and monthly returns we find that the decay factor for the daily data set is 0.94, and the decay factor for the monthly data set is 0.97.

5.4 Summary and concluding remarks

In this chapter we explained the methodology and practical issues surrounding the estimation of the RiskMetrics volatilities and correlations. Table 5.9 summarizes the important results about the RiskMetrics volatility and correlation forecasts.

Table 5.9
Summary of RiskMetrics volatility and correlation forecasts

Forecast	Expression*	Decay factor	# of daily returns used in production	Effective # of daily returns used in estimation†
1-day volatility	$\sigma_{1,t+1\mid t} = \sqrt{\lambda \sigma^2_{1,t\mid t-1} + (1-\lambda) r^2_{1,t}}$	0.94	550	75
1-day correlation	$\rho_{12,t+1\mid t} = \dfrac{\sigma^2_{12,t+1\mid t}}{\sigma_{1,t+1\mid t}\sigma_{2,t+1\mid t}}$	0.94	550	75
1-month volatility	$\sigma_{1,t+25\mid t} = 5 \cdot \sigma_{1,t+1\mid t}$	0.97	550	150
1-month correlation	$\rho_{12,t+25\mid t} = \rho_{12,t+1\mid t}$	0.97	550	150

* Note that in all calculations the sample mean of daily returns is set to zero.

† This number is a dependent of the decay factor explained in Section 5.3.1.2.

Lastly, recall from Chapter 4 that RiskMetrics assumes that returns are generated according to the model

[5.36] $$r_t = \sigma_t \varepsilon_t \qquad \varepsilon_t \to IID\ N(0, 1)$$

Now, given the recursive form of the EWMA model, a more complete version of the RiskMetrics model for any individual time series is

[5.37] $$\begin{aligned} r_t &= \sigma_t \varepsilon_t \qquad \varepsilon_t \to IID\ N(0, 1) \\ \sigma_t^2 &= \lambda \sigma_{t-1}^2 + (1-\lambda) r_{t-1}^2 \end{aligned}$$

Since Eq. [5.37] describes a process by which returns are generated, we can determine whether this model (evaluated at the optimal decay factor) can replicate the distinctive features of the observed data as presented in Chapter 4. We do so by generating a time series of daily returns from Eq. [5.37] for a given value of λ. A simulated time series from Eq. [5.37] with λ = 0.94 is shown in Chart 5.12.

Chart 5.12
Simulated returns from RiskMetrics model

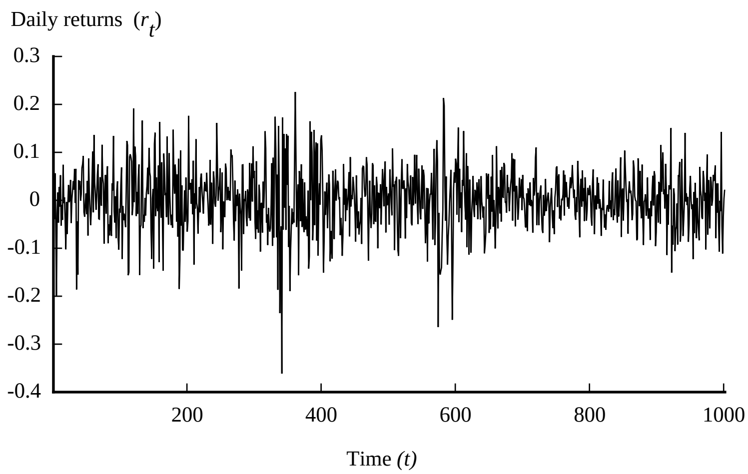

Chart 5.12 shows that the RiskMetrics model can replicate the volatility clustering feature noted in Chapter 4 (compare Chart 5.12 to Charts 4.6 and 4.7).