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

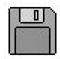

This diskette is located inside the back cover. It contains an Excel workbook that includes some of the examples shown in this document. Such examples are identified by the icon shown here.

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 I Risk Measurement Framework

Part II Statistics of Financial Market Returns

Part III Risk Modeling of Financial Instruments

Part IV RiskMetrics Data Sets

Part V Backtesting

Appendices

Reference

List of charts

List of tables

Part III

Risk Modeling of Financial Instruments

Chapter 6. Market risk methodology

Chapter 6. Market risk methodology

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

This chapter explains the methodology RiskMetrics uses to calculate VaR for portfolios that include multiple instruments such as simple bonds, swaps, foreign exchange, equity and other positions.

The chapter is organized as follows:

- Section 6.1 describes how to decompose various positions into cash flows.
- Section 6.2 covers how to convert or map the actual cash flows onto the corresponding RiskMetrics vertices.
- Section 6.3 explains two analytical approaches to measuring VaR.
- Section 6.4 presents a number of examples to illustrate the application of the RiskMetrics methodology.

6.1 Step 1—Identifying exposures and cash flows

The RiskMetrics building block for describing any position is a cash flow. A cash flow is defined by an amount of a currency, a payment date and the credit standing of the payor.

Once determined, these cash flows are marked-to-market. Marking-to-market a position's cash flows means determining the present value of the cash flows given **current** market rates and prices. This procedure requires current market rates, including the current on-the-run yield curve for newly issued debt, and a zero-coupon yield curve on instruments that pay no cash flow until maturity.[1] The zero coupon rate is the relevant rate for discounting cash flows received in a particular future period.[2]

We now describe how to express positions in fixed income, foreign exchange, equity, and commodities in terms of cash flows. The general process of describing a position in terms of cash flows is known as mapping.

6.1.1 Fixed Income

Interest rate positions describe the distribution of cash flows over time. Practitioners have applied various methods to express, or map, the cash flows of interest rate positions, the most common three being (1) duration map, (2) principal map, and (3) cash flow map. In this book we use the cash flow map method, but for comparison, present the two other methods as viable alternatives.

- Duration map

  The first and most common method to characterize a position's cash flows is by its duration (the weighted average life of a position's interest and principal payments). Macaulay duration is a measure of the weighted average maturity of an instrument's cash flows. Modified duration is a measure of a bond's price sensitivity to changes in interest rates. In general, duration provides risk managers with a simplified view of a portfolio's market risk. Its main drawback is that it assumes a linear relationship between price changes and yield changes. Moreover,

[1] See *The J. P. Morgan/Arthur Andersen Guide to Corporate Exposure Management* (p. 54, 1994).

[2] It is often suggested that implied forward rates are required to estimate the floating rates to be paid in future periods. In this document, however, we will show why forward rates are not necessarily required.

this approach works well when there are so-called parallel shifts in the yield curve but poorly when yield curves twist. Duration maps are used extensively in fixed income investment management. Many investment managers' activities are constrained by risk limits expressed in terms of portfolio duration.

- Principal map

  A second method, used extensively over the last two decades by commercial banks, is to describe a global position in terms of when principal payments occur. These "principal" maps form the basis for asset/liability management. ARBLs (Assets Repricing Before Liabilities) are used by banks to quantify interest rate risk in terms of cumulative assets maturing before liabilities. This method is employed most often when risks are expressed and earnings are accounted for on an accrual basis. The main problem with principal maps is that they assume that all interest payments occur at current market rates. This is often not a good assumption particularly when positions include fixed rate instruments with long maturities and when interest rates are volatile. Principal maps describe an instrument only as a function of the value and timing of redemption.

- Cash flow map

  The third method, and the one RiskMetrics applies is known as cash flow mapping. Fixed income securities can be easily represented as cash flows given their standard future stream of payments. In practice, this is equivalent to decomposing a bond into a stream of zero-coupon instruments. Complications in applying this technique can arise, however, when some of these cash flows are uncertain, as with callable or puttable bonds.

The following example shows how each of the mapping methodologies can be applied in practice. Chart 6.1 shows how a 10-year French OAT (FRF 100,000 francs nominal, 7.5% of April 2005) can be mapped under the approaches listed above:

- The duration map associates the market value of the instrument against the bond's Macaulay duration of 6.88 years.

- The principal map allocates the market value of the bond to the 10-year maturity vertex.

- The cash flow map shows the distribution over time of the current market value of all future streams (coupons + principals).

  As shown in Chart 6.1, the cash flow map (present valued) treats all cash flows separately and does not group them together as do the duration and principal maps. Cash flow mapping is the preferred alternative because it treats cash flows as being distinct and separate, enabling us to model the risk of the fixed income position better than if the cash flows were simply represented by a grouped cash flow as in the duration and principal maps.

Chart 6.1
French franc 10-year benchmark maps
amounts in thousands of market value

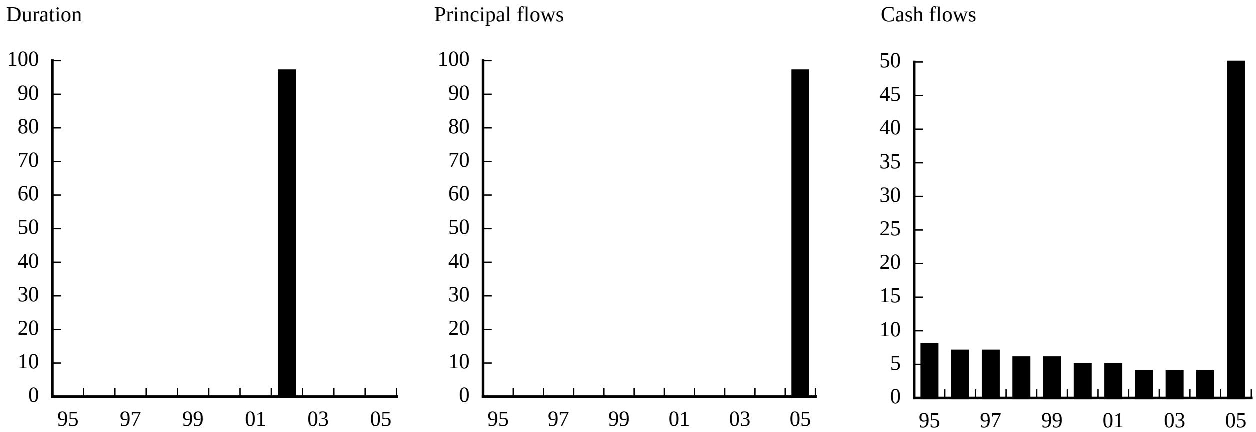

6.1.1.1 Simple bonds

Consider a hypothetical bond with a par value of 100, a maturity of 4 years and a coupon rate of 5%. Assume that the bond is purchased at time 0 and that coupon payments are payed on an annual basis at the beginning of each year. Chart 6.2 shows the bond's cash flows.[3]

In general, arrows pointing upwards signify cash inflows and arrows pointing downwards represent outflows. Also, a cash flow's magnitude is proportional to the length of the arrow; the taller (shorter) the arrow the greater (lower) the cash flow.

Chart 6.2
Cash flow representation of a simple bond

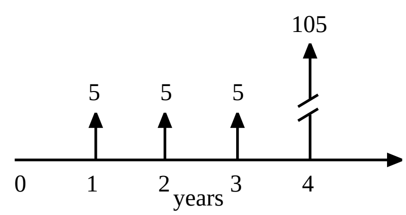

We can represent the cash flows of the simple bond in our example as cash flows from four zero-coupon bonds with maturities of 1,2,3 and 4 years. This implies that on a risk basis, there is no difference between holding the simple bond or the corresponding four zero-coupon bonds.

6.1.1.2 Floating rate notes (FRN)

A **floating rate note (FRN)** is an instrument that is based on a principal, P, that pays floating coupons. A FRN's coupon payment is defined as the product of the principal and a floating rate that is set some time in advance of the actual coupon payment. For example, if coupon payments are paid on a semiannual basis, the 6-month LIBOR rate would be used to determine the payment in 6 month's time. The coupon payments would adjust accordingly depending on the current 6-month LIBOR rate when the floating rate is reset. The principal is exchanged at both the beginning and end of the FRN's life.

3 We ignore the payment for the bond. That is, we do not account for the initial (negative) cash flow at time 0.

Chart 6.3 shows the cash flows for a hypothetical FRN lasting 4 years. The floating payments are represented by the gray shaded arrows. The black arrows represent fixed payments. All payments are assumed to occur on a yearly basis.

Chart 6.3
Cash flow representation of a FRN

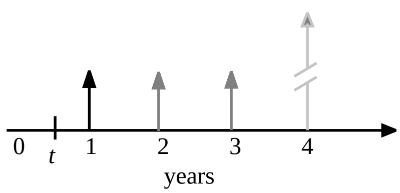

Notice that the first payment (at year 1) is known, and therefore, fixed. Also, the last payment represents the fact that the principal is known at the fourth year, but the final coupon payment is unknown. We now show how to evaluate the future floating payments.

Suppose that at time t (between 0 and 1 year), a risk manager is interested in analyzing the floating payment that will be received in year 3. The rate that determines this value is set in the second year and lasts one year. Now, implied forward rates are often used to forecast floating rates. The fundamental arbitrage relationship between current and future rates implies that the 1-year rate, as of year 2 satisfies the expression

[6.1] $$(1 + r_{2-t,t}) \cdot (1 + f_{1,2}) = (1 + r_{3-t,t})$$

where $r_{i,j}$ is the i-year rate set at time j and $f_{i,j}$ is the i period forward rate set at time j. So, for example, $f_{1,2}$ is the 1-year rate, beginning at the second year. It follows that the cash flow implied by this rate occurs in year 3. Since we know at time t both $r_{2-t,t}$ (the 2-t year rate) and $r_{3-t,t}$ (the 3-t year rate), we can solve for the implied forward rate as a function of observed rates. i.e.,

[6.2] $$f_{1,2} = \frac{(1 + r_{3-t,t})}{(1 + r_{2-t,t})} - 1$$

We can apply same technique to all other implied forward rates so that we can solve for $f_{1,1}$, $f_{1,2}$, $f_{1,3}$ and determine the expected future payments. The forecast coupon payment, for example, at time 3 is $P \cdot f_{1,2}$. The present value of this payment at time t, is simply $(P \cdot f_{1,2}) / (1 + r_{3-t,t})$. Substituting Eq. [6.2] into the expression for the discounted coupon payment yields,

[6.3] $$\frac{P \cdot f_{1,2}}{(1 + r_{3-t,t})} = \frac{P}{(1 + r_{2-t,t})} - \frac{P}{(1 + r_{3-t,t})}$$

Equation [6.3] shows that the expected coupon payment can be written in terms of known zero coupon rates. We can apply similar methods to the other coupon payments so that we can write the cash flows of the FRN as

[6.4] $$P_{FRN} = \frac{P \cdot r_{1,0}}{(1 + r_{1-t,t})} + \frac{P \cdot f_{1,1}}{(1 + r_{2-t,t})} + \frac{P \cdot f_{1,2}}{(1 + r_{3-t,t})} + \frac{P \cdot f_{1,3}}{(1 + r_{4-t,t})} + \frac{P}{(1 + r_{4-t,t})}$$

The right-hand side of Eq. [6.4] is equal to

[6.5]
$$\frac{P \cdot r_{1,0}}{(1+r_{1-t,t})} + \left(\frac{P}{(1+r_{1-t,t})} - \frac{P}{(1+r_{2-t,t})}\right) + \left(\frac{P}{(1+r_{2-t,t})} - \frac{P}{(1+r_{3-t,t})}\right) + \left(\frac{P}{(1+r_{3-t,t})} - \frac{P}{(1+r_{4-t,t})}\right) + \frac{P}{(1+r_{4-t,t})}$$

Equation [6.5] collapses to the present value

[6.6]
$$\frac{P \cdot (1+r_{1,0})}{(1+r_{1-t,t})}$$

Chart 6.4 shows that the cash flow of the FRN from the time t perspective, is $P(1+r_{1,0})$. Therefore, we would treat the FRN's cash flows as a cash flow from a zero coupon bond with maturity 1-*t* period.

Chart 6.4
Estimated cash flows of a FRN

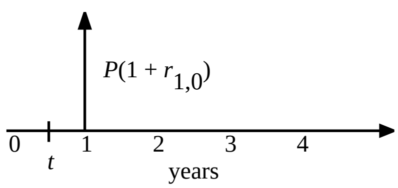

Notice that if the cash flows in Chart 6.3 were computed relative to time zero (the start of the FRN), rather than to time *t*, the cash flow would be simply *P* at *t* = 0, representing the par value of the FRN at its start.

6.1.1.3 Simple interest-rate swaps
Investors enter into interest-rate swaps to change their exposure to interest rate uncertainty by exchanging interest flows. In order to understand how to identify a simple interest-rate swap's cash flows, a swap should be thought of as a portfolio consisting of one fixed and one floating rate instrument. Specifically, the fixed leg is represented by a simple bond without an exchange of principal. The floating leg is a FRN with the caveat that the principal is used only to determine coupon payments, and is not exchanged.

Chart 6.5 shows the cash flows of an interest-rate swap that receives fixed rate and pays the floating rate.

Chart 6.5
Cash flow representation of simple interest rate swap

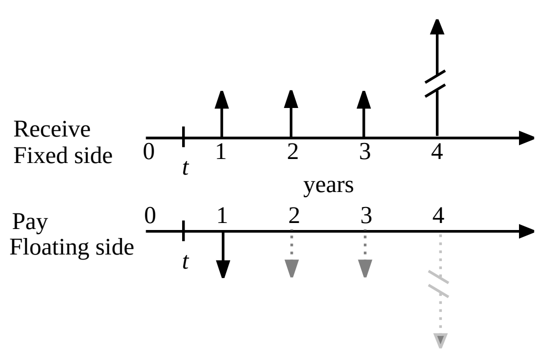

We compute the cash flows relative to time t, (again, between 0 and 1 year) after the start of the swap. The cash flows on the fixed side are simply the fixed coupon payments over the next 4 years which, as already explained in Section 6.1.1.1, are treated as holding four zero-coupon bonds. The cash flows on the floating side are derived in the exact manner as the payments for the FRN (except now we are short the floating payments). The present value of the cash flow map of the floating side of the swap is given by Eq. [6.7]

$$[6.7] \qquad -\frac{P \cdot (1 + r_{1,0})}{(1 + r_{1-t,t})},$$

where P is the principal of the swap. Notice the similarity between this cash flow and that given by Eq. [6.6] for the FRN. Hence, we can represent the cash flows on the floating side of the swap as being short a zero coupon bond with maturity 1-t.

6.1.1.4 Forward starting swap

A forward starting swap is an instrument where one enters into an agreement to swap interest payments at some future date. Unlike a simple swap none of the floating rates are fixed in advance. Chart 6.6 shows the cash flows of a forward starting swap.

Chart 6.6

Cash flow representation of forward starting swap

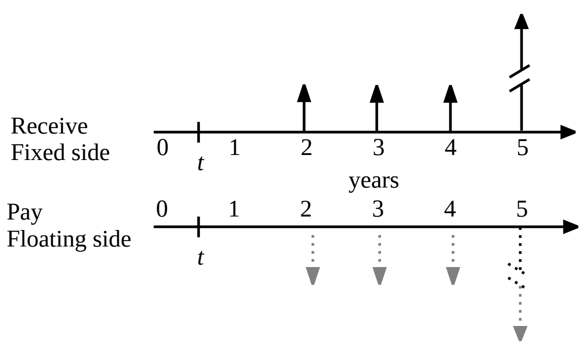

Suppose that an investor enters into a forward starting swap with 5 years to maturity at some time t (the trade date), and the start date of the swap, (i.e., the date when the floating rates are fixed) is year 2. Starting in year 3, payments are made every year until year 5. The cash flows for this instrument are essentially the same as a simple interest-rate swap, but now the first floating payment is unknown.

The cash flows on the fixed side are simply the cash flows discounted back to time t. On the floating side, the cash flows are, again, determined by the implied forward rates. The cash flow map for the (short) floating payments is represented by Eq. [6.8].

$$[6.8] \qquad -\frac{P}{(1 + r_{2-t,t})}$$

Chart 6.7 depicts this cash flow.

Chart 6.7
Cash flows of the floating payments in a forward starting swap

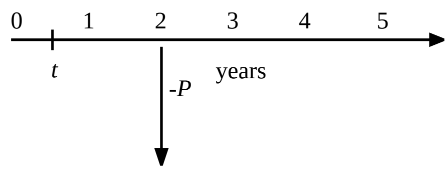

Notice that this cash flow map is equivalent to being short a 2-*t* zero coupon bond.

6.1.1.5 Forward rate agreement (FRA)
A **forward rate agreement (FRA)** is an interest rate contract. It locks in an interest rate, either a borrowing rate (buying a FRA) or a lending rate (selling a FRA) for a specific period in the future. FRAs are similar to futures but are over-the-counter instruments and can be customized for any maturity.

A FRA is a notional contract. Therefore, there is no exchange of principal at the expiry date (i.e., the fixing date). In effect, FRAs allow market participants to lock in a forward rate that equals the implied break even rate between money market and term deposits.[4] To understand how to map the cash flows of a FRA, let's consider a simple, hypothetical example of a purchase of a 3 vs. 6 FRA at r% on a notional amount P. This is equivalent to locking in a borrowing rate for 3 months starting in 3 months. The notation 3 vs. 6 thus refers to the start date of the underlying versus the end date of the underlying, with the start date being the delivery date of the contract.

Chart 6.8 depicts the cash flows of this FRA.

Chart 6.8
Cash flow representation of FRA

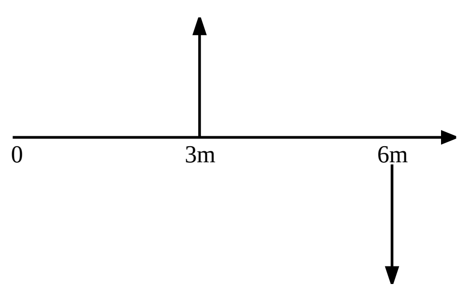

We can replicate these cash flows by going long the current 3-month rate and short the 6-month rate as shown in Chart 6.9.

[4] For more details on FRAs, refer to Valuing and Using FRAs (Hakim Mamoni, October, 1994, JP Morgan publication).

Chart 6.9
Replicating cash flows of 3-month vs. 6-month FRA

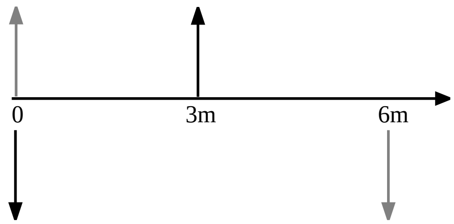

Note that the gray arrows no longer represent floating payments. The gray and black arrows represent the cash flows associated with going short a 6-month zero coupon bond and long a 3-month zero coupon bond, respectively. The benefit of working with the cash flows in Chart 6.9 rather than in Chart 6.8, is that the latter requires information on forward rates whereas the former does not.

6.1.1.6 Interest rate future
We now consider the cash flow map of a 3-month Eurodollar future contract that expires in 3 months'. Taking time 0 to represent the current date, we represent the future's cash flows by an outflow in 3 months and an inflow in 6 months, as shown in Chart 6.10.

Chart 6.10
Cash flow representation of 3-month Eurodollar future

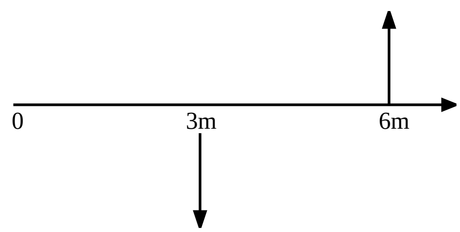

To be more specific, if the current USD 3-month Eurodollar deposit rate is 7.20%, a purchaser of this futures contract would face a cash outflow of USD 981,800 in 3 months and a cash inflow of USD 1,000,000 in 6 months. We can then represent these cash flows as being short the current 3-month rate and investing this money in the current 6-month rate. Hence, the cash flows of this Eurodollar futures contract can be replicated by a short 3-month position and a long 6-month position as shown in Chart 6.11.

Chart 6.11
Replicating cash flows of a Eurodollar futures contract

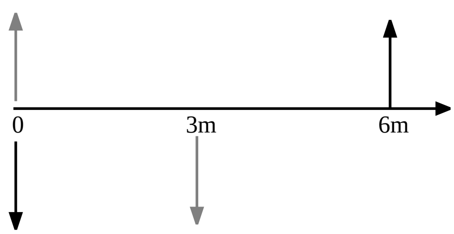

6.1.2 Foreign exchange

Financial positions are described in terms of a base or "home" currency. For example, American institutions report risks in U.S. dollars, while German institutions use Deutsche marks. A risk manager's risk profile is not independent of the currency in which risk is reported. For example, consider two investors. One investor is based in US dollars, the other in Italian lira. Both investors purchase an Italian government bond. Whereas the USD based investor is exposed to both interest rate and exchange rate risk (by way of the ITL/USD exchange rate), the lira based investor is exposed only to interest rate risk. Therefore, an important step to measure foreign exchange risk is to understand how cash flows are generated by foreign exchange positions.

6.1.2.1 Spot positions

Describing cash flows of spot foreign exchange positions is trivial. Graphically, up and down arrows represent long and short positions in foreign exchange, respectively.

6.1.2.2 Forward foreign exchange positions

A foreign exchange (FX) forward is an agreement to exchange at a future date, an amount of one currency for another at a specified forward rate. Mapping a forward foreign exchange position is facilitated by the ability to express the forward as a function of two interest rates and a spot foreign exchange rate.[5] For example, Chart 6.12 shows the cash flows of an FX forward that allows an investor to buy Deutsche marks with US dollars in 6 months' time at a prespecified forward rate.

Chart 6.12
FX forward to buy Deutsche marks with US dollars in 6 months

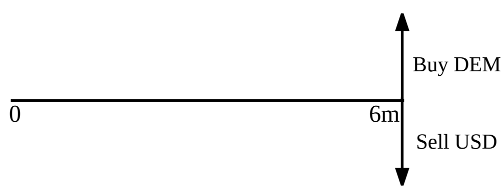

We can replicate these cash flows by borrowing dollars at time 0 at the 6-month interest rate $r_{USD, t}$ and immediately investing these dollars in Germany at a rate $r_{DEM, t}$, This scenario would generate the cash flows which, at the 6-month mark, are identical to those of the forward contract. These cash flows are shown in Chart 6.13.

Chart 6.13
Replicating cash flows of an FX forward

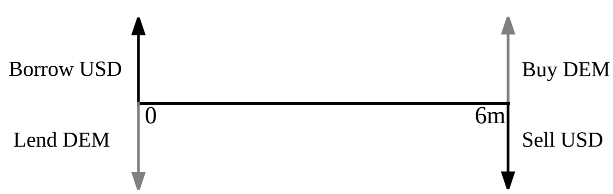

The ability to replicate future foreign exchange cash flows with interest rate positions results from what is known as interest rate parity (IRP). We now demonstrate this condition. Let the spot rate, S_t, of the home currency expressed in units of foreign currency, (e.g., if the home currency is the US dollars and the foreign currency is Deutsche marks, S_t is expressed in US dollars per Deutsche

[5] For simplicity, we ignore other factors such as transaction costs and possible risk premia.

marks (USD/DEM)). The forward rate, F_t, is the exchange rate observed at time t, which guarantees a spot rate at some future time T. Under interest rate parity the following condition holds

$$F_t = S_t \frac{(1 + r_{USD,t})}{(1 + r_{DEM,t})} \qquad [6.9]$$

It follows from IRP that the ability to convert cash flows of an FX forward into equivalent borrowing and lending positions implies that holding an FX forward involves cash flows that are exposed to both foreign exchange and interest risk.

6.1.2.3 Currency swaps

Currency swaps are swaps for which the two legs of the swap are each denominated in a different currency. For example, one party might receive fixed rate Deutsche marks, the other floating rate US dollars. Unlike an interest-rate swap, the notional principal in a currency swap is exchanged at the beginning and end of the swap.[6]

Chart 6.14 shows the cash flows for a hypothetical currency swap with a maturity of 4 years and paying fixed rate Deutsche marks and floating rate US dollars on an annual basis. For completeness, we present the cash flows associated with the initial exchange of principal.

Chart 6.14

Actual cash flows of currency swap

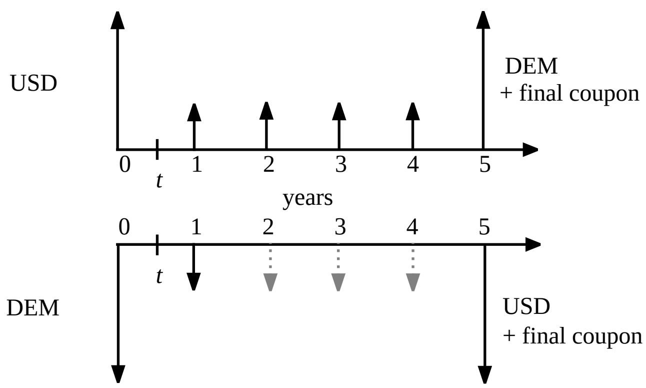

From the perspective of holding the swap at time t between 0 and year 1, the fixed leg of the swap has the same cash flows as the simple bond presented in Section 6.1.1.1. The cash flows of the floating leg are the same as that as a short position in a FRN.

6.1.3 Equities

The cash flows of equity are simple spot positions expressed in home currency equivalents. Equity positions held in foreign countries are subject to foreign exchange risk in addition to the risk from holding equity.

[6] There are currency swaps where one or both of the notional amounts are not exchanged.

6.1.4 Commodities

Exposures to commodities can be explained using a framework similar to that of interest rates. Risks arise in both the spot market (you purchase a product today and store it over time) and from transactions that take place in the future (e.g., physical delivery of a product in one month's time).

6.1.4.1 Commodity futures contract

Commodity futures contracts enable investors to trade products for future delivery with relative ease and also serve as a price setting and risk transferring mechanisms for commodity producers. These contracts provide market participants with valuable information about the term structure of commodities prices.

6.1.4.2 Commodity swap

Institutions do not have to limit themselves to futures contracts when they participate in the commodity markets. They can enter into swaps to change their exposure to interest rates, currency, and/or commodity risks. A typical commodity swap entails an institution to paying (receiving) fixed amounts in exchange for receiving (paying) variable amounts with respect to an index (e.g., an average of the daily price of the nearby natural gas futures contract).

In many respects, commodity swaps are similar to interest-rate swaps. Unlike an interest-rate swap the underlying instrument of a commodity swap can be of variable quality thereby making the terms of the transaction more complex.

6.2 Step 2—Mapping cash flows onto RiskMetrics vertices

In the last section we described cash flows generated by particular classes of instruments. Financial instruments, in general, can generate numerous cash flows, each one occurring at a unique time. This gives rise to an unwieldy number of combinations of cash flow dates when many instruments are considered. As a result, we are faced with the impractical task of having to compute an intractable number of volatilities and correlations for the VaR calculation. To more easily estimate the risks associated with instruments' cash flows, we need to simplify the time structure of these cash flows.

The RiskMetrics method of simplifying time structure involves cash flow mapping, i.e., redistributing (mapping) the observed cash flows onto so-called RiskMetrics vertices, to produce RiskMetrics cash flows.

6.2.1 RiskMetrics vertices

All RiskMetrics cash flows use one or more of the 14 RiskMetrics vertices shown below (and on page 107).

1m 3m 6m 12m 2yr 3yr 4yr 5yr 7yr 9yr 10yr 15yr 20yr 30yr

These vertices have two important properties:

- They are fixed and hold at any time now and in the future for all instruments, linear and non-linear. (J.P. Morgan can occasionally redefine these vertices to keep up with market trends.)

- RiskMetrics data sets provide volatilities and correlations for each of these vertices (and only for these vertices).

Mapping an actual cash flow involves splitting it between the two closest RiskMetrics vertices (unless the cash flow happens to coincide with a RiskMetrics vertex). For example, a cash flow occurring in 6 years is represented as a combination of a 5-year and a 7-year cash flow. Chart 6.15

shows how the actual cash flow occurring at year 6 is split into the synthetic (RiskMetrics) cash flows occurring at the 5- and 7-year vertices.

Chart 6.15
RiskMetrics cash flow mapping

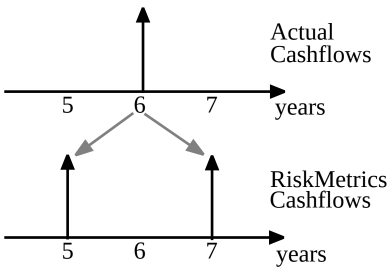

The two fractions of the cash flow are weighted such that the following three conditions hold:

1. **Market value is preserved.** The total market value of the two RiskMetrics cash flows must be equal to the market value of the original cash flow.
2. **Market risk is preserved.** The market risk of the portfolio of the RiskMetrics cash flows must also be equal to the market risk of the original cash flow.
3. **Sign is preserved.** The RiskMetrics cash flows have the same sign as the original cash flow.

In the trivial case that the actual vertex and RiskMetrics vertex coincide, 100% of the actual cash flow is allocated to the RiskMetrics vertex.

It is important to understand that RiskMetrics cash flow mapping differs from conventional mapping methods in the three conditions that it stipulates. A common practice used to date throughout the financial industry has been to follow **two** standard rules when allocating cash flows between vertices:

1. **Maintain present value.** For example, the sum of the cash flows maturing in 5 and 7 years must be equal to the original cash flow occurring in year 6.
2. **Maintain duration.** The duration of the combination of 5- and 7-year cash flows must also be equal to the duration of the 6-year cash flow.

Cash flow maps like these are similar to a barbell type trade, where an existing position is replaced by a combination of two instruments distributed along the yield curve under the condition that the trade remains duration neutral. Barbell trades are entered into by investors who are duration-constrained but have a view on a shift in the yield curve. What is a perfectly defensible investment strategy, however, cannot be simply applied to risk estimation.

6.2.2 Computing RiskMetrics cash flows

For allocating actual cash flows to RiskMetrics vertices, RiskMetrics proposes a methodology that is based on the variance (σ^2) of financial returns. The advantage of working with the variance is that it is a risk measure closely associated with one of the ways RiskMetrics computes VaR, namely the simple VaR method as opposed to the delta-gamma or Monte Carlo methods.

In order to facilitate the necessary mapping, the RiskMetrics data sets provide users with volatilities on, and correlations across many instruments in 33 markets. For example, in the US government bond market, RiskMetrics data sets provide volatilities and correlations on the 2-, 3-, 4-, 5-, 7-, 9-, 10-, 15-, 20-, and 30-year zero coupon bonds.

We now demonstrate how to convert actual cash flows to RiskMetrics cash flows, continuing with the example of allocating a cash flow in year 6 to the 5- and 7-year vertices (Chart 6.15). We denote the allocations to the 5- and 7-year vertices by α and $(1-\alpha)$, respectively. The procedure presented below is not restricted to fixed income instruments, but applies to all future cash flows.

1. Calculate the actual cash flow's interpolated yield:

   We obtain the 6-year yield, y_6, from a linear interpolation of the 5- and 7-year yields provided in the RiskMetrics data sets. Using the following equation,

   [6.10] $$y_6 = \hat{a} y_5 + (1 - \hat{a}) y_7 \qquad 0 \le \hat{a} \le 1$$

   where y_6 = interpolated 6-year zero yield
   $\hat{a}$ = linear weighting coefficient, $\hat{a} = 0.5$ in this example
   y_5 = 5-year zero yield
   y_7 = 7-year zero yield

   If an actual cash flow vertex is not equidistant between the two RiskMetrics vertices, then the greater of the two values, $\hat{a}$ and $(1 - \hat{a})$, is assigned to the closer RiskMetrics vertex.

2. Determine the actual cash flow's present value:

   From the 6-year zero yield, y_6, we determine the present value, P_6, of the cash flow occurring at the 6-year vertex. (In general, P_i denotes the present value of a cash flow occurring in i years.)

3. Calculate the standard deviation of the price return on the actual cash flow:

   We obtain the standard deviation, σ_6, of the return on the 6-year zero coupon bond, by a linear interpolation of the standard deviations of the 5- and 7-year price returns, i.e., σ_5 and σ_7, respectively.

   Note that σ_5 and σ_7 are provided in the RiskMetrics data sets as the VaR statistics $1.65\sigma_5$ and $1.65\sigma_7$, respectively. Hence, $1.65\sigma_6$ is the interpolated VaR.

   To obtain σ_6, we use the following equation:

   [6.11] $$\sigma_6 = \hat{a}\sigma_5 + (1 - \hat{a})\sigma_7 \qquad 0 \le \hat{a} \le 1$$

   where

   $\hat{a}$ = linear weighting coefficient from Eq. [6.10]
   σ_5 = standard deviation of the 5-year return
   σ_7 = standard deviation of the 7-year return

4. Compute the allocation, α and (1-α), from the following equation:

[6.12]
$$\text{Variance}\,(r_{6\text{yr}}) = \text{Variance}\,[\alpha r_{5\text{yr}} + (1-\alpha)\,r_{7\text{yr}}],\ \text{or the equivalent}$$
$$\sigma_6^2 = \alpha^2\sigma_5^2 + 2\alpha\,(1-\alpha)\,\rho_{5,7}\sigma_5\sigma_7 + (1-\alpha)^2\sigma_7^2$$

where $\rho_{5,7}$, is the correlation between the 5- and 7- year returns. (Note that $\rho_{5,7}$ is provided in the correlation matrix in RiskMetrics data sets).

Equation [6.12] can be written in the quadratic form

[6.13]
$$a\alpha^2 + b\alpha + c = 0$$

where

$$a = \sigma_5^2 + \sigma_7^2 - 2\rho_{5,7}\sigma_5\sigma_7$$
$$b = 2\rho_{5,7}\sigma_5\sigma_7 - 2\sigma_7^2$$
$$c = \sigma_7^2 - \sigma_6^2$$

The solution to α is given by

[6.14]
$$\alpha = \frac{-b \pm \sqrt{b^2 - 4ac}}{2a}$$

Notice that Eq. [6.14] yields two solutions (roots). We choose the solution that satisfies the three conditions listed on page 118.

5. Distribute the actual cash flow onto the RiskMetrics vertices:

   Split the actual cash flow at year 6 into two components, α and (1-α), where you allocate α to the 5-year RiskMetrics vertex and (1-α) to the 7-year RiskMetrics vertex.

Using the steps above, we compute a RiskMetrics cash flow map from the following real-world data. Suppose that on July 31, 1996, the cash flow occurring in 6 years is USD 100. The RiskMetrics daily data sets provide the statistics shown in Table 6.1, from which we calculate the data shown in Table 6.2.[7]

[7] Recall that RiskMetrics provides VaR statistics—that is, 1.65 times the standard deviation.

Table 6.1
Data provided in the daily RiskMetrics data set

y_5	5-year yield	6.605%
y_7	7-year yield	6.745%
$1.65\sigma_5$	volatility on the 5-year bond price return	0.5770%
$1.65\sigma_7$	volatility on the 7-year bond price return	0.8095%
$\rho_{5,7}$	correlation between the 5- and 7-year bond returns	0.9975

Table 6.2
Data calculated from the daily RiskMetrics data set

y_6	6-year yield (from Eq. [6.10], where $\hat{a} = 0.5$)	6.675%
σ_6	standard deviation on the 6-year bond price return	0.4202%
σ_6^2	variance on the 6-year bond price return (from Eq. [6.11])	1.765×10^{-3}%
σ_5^2	variance on the 5-year bond price return	1.223×10^{-3}%
σ_7^2	variance on the 7-year bond price return	2.406×10^{-3}%

To solve for α, we substitute the values in Tables 6.1 and 6.2 into Eq. [6.13] to find the following values:

$$a = 2.14 \times 10^{-6}$$
$$b = -1.39 \times 10^{-5}$$
$$c = 6.41 \times 10^{-6}$$

which in Eq. [6.14], yields the solutions $\alpha = 5.999$ and $\alpha = 0.489$. We disqualify the first solution since $(1-\alpha)$ violates the sign preservation condition (page 118).

From Step 2 on page 119, the present value of USD 100 in 6 years is USD 93.74, i.e., the 6-year cash flow. We allocate 49.66% (USD 46.55) of it to the 5-year vertex and 50.33% (USD 47.17) to the 7-year vertex. We thus obtain the RiskMetrics cash flow map shown in Chart 6.15.

The preceding example demonstrated how to map a single cash flow to RiskMetrics vertices. In practice portfolios often contain many cash flows, each of which has to be mapped to the RiskMetrics vertices. In such instances, cash flow mapping simply requires a repeated application of the methodology explained in this section.

6.3 Step 3—Computing Value-at-Risk

This section explains two analytical approaches to measuring Value-at-Risk: simple VaR for linear instruments, and delta-gamma VaR for nonlinear instruments, where the terms "linear" and "nonlinear" describe the relationship of a position's underlying returns to the position's relative change in value. (For more information about simple VaR methodology, see Section 6.3.2. For more information about delta-gamma methodology, see Section 6.3.3.)

In the simple VaR approach, we assume that returns on securities follow a conditionally multivariate normal distribution (see Chapter 4) and that the relative change in a position's value is a linear function of the underlying return. Defining VaR as the 5th percentile of the distribution of a portfolio's relative changes, we compute VaR as 1.65 times the portfolio's standard deviation, where the multiple 1.65 is derived from the normal distribution. This standard deviation depends on the volatilities and correlations of the underlying returns and on the present value of cash flows.

In the delta-gamma approach, we continue to assume that returns on securities are normally distributed, but allow for a nonlinear relationship between the position's value and the underlying returns. Specifically, we allow for a second-order or gamma effect, which implies that the distribution of the portfolio's relative change is no longer normal. Therefore, we cannot define VaR as 1.65 times the portfolio's standard deviation. Instead, we compute VaR in two basic steps. First, we calculate the first four moments of the portfolio's return distribution, i.e., the mean, standard deviation, skewness and kurtosis. Second, we find a distribution that has the same four moments and then calculate the fifth percentile of this distribution, from which we finally compute the VaR.

The choice of approach depends on the type of positions that are at risk, i.e., linear or non-linear positions, as defined above.

6.3.1 Relating changes in position values to underlying returns

This section explains the linearity and nonlinearity of instruments in the context of RiskMetrics methodology.

Value-at-Risk measures the market risk of a portfolio. We define a portfolio as a set of positions, each of which is composed of some underlying security. In order to calculate the risk of the portfolio, we must be able to compute the risks of the positions that compose the portfolio. This requires an understanding of how a position's value changes when the value on its underlying security changes. Thus, we classify positions into simple positions, which are linear, and into derivative positions, which can be further broken down into linear and nonlinear derivative positions.

As an example of a simple position, the relative change in value of a USD 100 million dollar position in DEM is a linear function of the relative change in value in the USD/DEM exchange rate (i.e., the return on the USD/DEM exchange rate).

The value of derivative positions depends on the value of some other security. For example, the value of a forward rate agreement, a linear derivative, depends on the value of some future interest rate. In contrast, other derivatives may have a nonlinear relationship between the relative change in value of the derivative position and the value of the underlying security. For example, the relative change in value of an option on the USD/FRF exchange rate is a nonlinear function of the return on that rate.

Chart 6.16 shows how the return on a position varies with the return on the underlying security.

Chart 6.16
Linear and nonlinear payoff functions

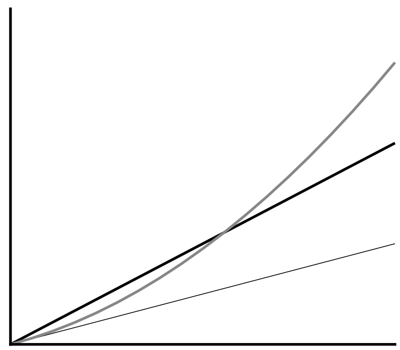

The straight lines in Chart 6.16 signify a constant relationship between the position and underlying security. The black line represents a one-to-one relationship between the position value and the underlying security. Note that for a payoff to be linear, the movement between the position value and the underlying security's value does not have to be one-to-one. For example, the change in value of a simple option can be expressed in terms of the "delta" (slope) of the underlying security, where the delta varies between −1 and +1. Chart 6.16 shows a payoff function where delta is 0.5 (gray line).

When payoffs are nonlinear there is no longer a "straight line" relationship between the position value and the underlying security's value. Chart 6.16 shows that the payoff line is curved such that the position value can change dramatically as the underlying security value increases. The convexity of the line is quantified by the parameter "gamma".

In summary, linear payoffs are characterized by a constant slope, delta. Their convexity, gamma, is always equal to zero. VaR for such instruments is calculated from the simple VaR methodology (Section 6.3.2). For nonlinear payoffs, delta can take on any value between –1 and +1, while gamma is always non-zero, accounting for the observed curvature of the payoff function. Nonlinear instruments are thus treated by the delta-gamma methodology (although the same methodology can also be used to handle linear instruments. See Section 6.3.3 on page 129).

Table 6.3 lists selected positions (instruments), their underlying returns, and the relationship between the two.

Table 6.3
Relationship between instrument and underlying price/rate

Type of position	Instrument*	Underlying price/rate†
Simple (linear):	Bond	Bond price§
	Stock	Local market index
	Foreign exchange	FX rate
	Commodity	Commodity price
	IR swap	Swap price
Linear derivative:	Floating rate note	Money market price

Table 6.3 (continued)
Relationship between instrument and underlying price/rate

Type of position	Instrument*	Underlying price/rate†
	FX forward	FX rate/money market price
	Forward rate agreement	Money market price
	Currency swap	Swap price/FX rate
Nonlinear derivative:	Stock Option	Stock price
	Bond Option	Bond price
	FX Option	FX rate

* Treated by $\tilde{r}_t$. See the remainder of Section 6.3.
† Treated by r_t. See the remainder of Section 6.3.
§ Note, however, the relationship between a bond price and its yield is nonlinear.

6.3.1.1 Linear positions
Using the qualitative information in the preceding section, we now formally derive the relationship between the relative change in the value of a position and an underlying return for linear instruments.

We denote the relative change in value of the *i*th position, at time *t*, as $\tilde{r}_{i,t}$. In the simple case where there is a linear one-to-one correspondence between the relative change in value of this position and some underlying return $r_{j,t}$, we have $\tilde{r}_{i,t} = r_{j,t}$.[8] In general, we denote a position that is linearly related to an underlying return as $\tilde{r}_{i,t} = \delta r_{j,t}$, where δ is a scalar.

Notice that in the case of fixed income instruments, the underlying value is defined in terms of prices on zero equivalent bonds (Table 6.3). Alternatively, underlying returns could have been defined in terms of yields. For example, in the case of bonds, there is no longer a one-to-one correspondence between a change in the underlying yield and the change in the price of the instrument. In fact, the relationship between the change in price of the bond and yield is nonlinear. Since we only deal with zero-coupon bonds we focus on these. Further, we work with continuous compounding.

Assuming continuous compounding, the price of an N-period zero-coupon bond at time *t*, P_t, with yield y_t is

$$[6.15] \qquad P_t = e^{-y_t N}$$

A second order approximation to the relative change in P_t yields

$$[6.16] \qquad \tilde{r}_t = -y_t N (\Delta y_t / y_t) + \frac{1}{2} (y_t N)^2 (\Delta y_t / y_t)^2$$

Now, if we define the return r_t in terms of relative yield changes, i.e., $r_t = (\Delta y_t / y_t)$, then we have

$$[6.17] \qquad \tilde{r}_t = -y_t N (r_t) + \frac{1}{2} (y_t N)^2 (r_t)^2$$

[8] Technically, this results from the fact that the derivative of the price of the security with respect to the underlying price is 1.

Equation [6.17] reveals two properties: (1) If we ignore the second term on the right-hand side we find that the relative price change is linearly, but not in one-to-one correspondence, related to the return on yield. (2) If we include the second term, then there is a nonlinear relationship between return, r_t, and relative price change.

6.3.1.2 Nonlinear positions (options)

In options positions there is a nonlinear relationship between the change in value of the position and the underlying return. We explain this relationship with a simple stock option. For a given set of parameters denote the option's price by $V(P_t, K, \tau, \rho, \sigma)$ where P_t is the spot price on the underlying stock at time t, K is the option's exercise price, τ is the time to maturity of the option in terms of a year, ρ is the riskless rate of a security that matures when the option does, and σ is the standard deviation of the log stock price change over the horizon of the option.

In order to obtain an expression for the return on the option, $\tilde{r}_{i,t}$, we approximate the future value of the option $V(P_{t+1}, K, \tau, \rho, \sigma)$ with a second-order Taylor series expansion around the current values (spot rates), $P_t, K, \tau, \rho,$ and σ. This yields,

$$[6.18] \qquad V(P_{t+1}, K, \tau, \rho, \sigma) \cong V_0(P_t, K, \tau, \rho, \sigma) + \left(\frac{\partial V}{\partial P}\right)(P_{t+1} - P_t) + \frac{1}{2}\left(\frac{\partial^2 V}{\partial P^2}\right)(P_{t+1} - P_t)^2$$

which can be rewritten in the more concise form

$$[6.19] \qquad dV \cong \delta(dP) + \frac{1}{2}\Gamma(dP)^2$$

Notice that dV, the change in value of the option, is in units of price P thus δ is unitless and Γ is in units of $1/P$. Writing Eq. [6.19] in terms of relative changes, we get

$$[6.20] \qquad \tilde{r}_t \cong \eta \cdot \left[\delta r_t + \frac{1}{2}\Gamma P_t r_t^2\right]$$

where η measures the leverage effect of holding the option, δ measures the relative change in the value of the option given a change in the value of the price P_t, Γ measures the relative change in the value of the option given a change in the value of δ.

As Eq. [6.20] shows, the relative change, $\tilde{r}_t$, in the option is a nonlinear function of r_t, the return on the underlying stock price, since it involves the term r_t^2.

6.3.2 Simple VaR calculation

In this section we provide the general formula to compute VaR for linear instruments. (These instruments include the first nine listed in Table 6.3.) The example provided below deals exclusively with the VaR calculation at the 95% confidence interval using the data provided by RiskMetrics.

Consider a portfolio that consists of N positions and that each of the positions consists of one cash flow on which we have volatility and correlation forecasts. Denote the relative change in value of the nth position by $\tilde{r}_{n,t}$. We can write the change in value of the portfolio, $\tilde{r}_{p,t}$, as

$$[6.21] \qquad \tilde{r}_{p,t} = \sum_{n=1}^{N} \omega_n \tilde{r}_{n,t} = \sum_{n=1}^{N} \omega_n \delta_n r_{n,t}$$

where ω_n is the total (nominal) amount invested in the *n*th position. For example, suppose that the total current market value of a portfolio is $100 and that $10 is allocated to the first position. It follows that $\omega_1 = \$10$.

Now, suppose that the VaR forecast horizon is one day. In RiskMetrics, the VaR on a portfolio of simple linear instruments can be computed by 1.65 times the standard deviation of $\tilde{r}_{p,t}$—the portfolio return, one day ahead.The expression of VaR is given as follows.

[6.22] $$VaR_t = \sqrt{\tilde{\sigma}_{t|t-1} R_{t|t-1} \tilde{\sigma}_{t|t-1}^T} \quad \text{(Value-at-Risk estimate)}$$

where

[6.23] $$\tilde{\sigma}_{t|t-1} = [1.65\sigma_{1,t|t-1}\omega_1\delta_1 \quad 1.65\sigma_{2,t|t-1}\omega_2\delta_2 \ \ldots \quad 1.65\sigma_{N,t|t-1}\omega_N\delta_N]$$

is the individual VaR vector (1xN) and

[6.24] $$R_{t|t-1} = \begin{bmatrix} 1 & \rho_{12,t|t-1} & \cdots & \rho_{1N,t|t-1} \\ \rho_{21,t|t-1} & 1 & \ldots & \ldots \\ \ldots & \ldots & \ldots & \ldots \\ \rho_{N1,t|t-1} & \ldots & \ldots & 1 \end{bmatrix}$$

is the NxN correlation matrix of the returns on the underlying cash flows.

The above computations are for portfolios whose returns can be reasonably approximated by the conditional normal distribution. In other words, it is assumed that the portfolio return follows a conditional normal distribution.

6.3.2.1 Fixed income instruments

In this section we address two important issues related to calculating the VaR on a portfolio of fixed income instruments. The first issue relates to what variable should be used to measure volatility and correlation. In other words, should we compute volatilities and correlations on prices or on yields? The second issue deals with incorporating the “roll down” and “pull-to-par” effects of bonds into VaR calculations.

We discussed in Section 6.3.1.1 that one may choose to model either the yield (interest rate) or the price of a fixed income instrument. **RiskMetrics computes the price volatilities and correlations on all fixed income instruments**. This is done by first computing zero rates for all instruments with a maturity of over one year, and then constructing prices from these series using the expression (continuous compounding).

[6.25] $$P_t = e^{-y_t N}$$

where y_t is the current yield on the N-period zero-coupon bond.

For money market rates, i.e., instruments with a maturity of less than one-year, prices are constructed from the formula

[6.26] $$P_t = \frac{1}{(1 + y_t)^N}$$

Since practitioners often think of volatilities on fixed income instruments in terms of yield, we present the price volatility in terms of yield volatility. Starting with Eq. [6.25], we find the price return to be

$$r_t = \ln(P_t / P_{t-1}) = N(y_{t-1} - y_t) \qquad [6.27]$$

Therefore, the standard deviation of price returns under continuous compounding is given by the expression

$$\sigma_t = N\sigma(y_{t-1} - y_t) \qquad [6.28]$$

where $\sigma(y_{t-1} - y_t)$ is the standard deviation of $y_{t-1} - y_t$. What Eq. [6.28] states is that price return volatility is the maturity on the underlying instrument times the standard deviation of the absolute changes in yields.

Performing the same exercise on Eq. [6.26] we find the price return to be

$$\begin{aligned} r_t &= \ln(P_t / P_{t-1}) \\ &= N\left[\ln\left(\frac{1 + y_{t-1}}{1 + y_t}\right)\right] \end{aligned} \qquad [6.29]$$

In this case (discrete compounding) the standard deviation of price returns is

$$\sigma_t = N\sigma\left[\ln\left(\frac{1 + y_t}{1 + y_{t-1}}\right)\right] \qquad [6.30]$$

where $\sigma\left[\ln\left(\frac{1 + y_t}{1 + y_{t-1}}\right)\right]$ is the standard deviation of $\ln\left(\frac{1 + y_t}{1 + y_{t-1}}\right)$.

We now explain how to incorporate the unique features of fixed income instruments in VaR calculations.[9] Traditionally, RiskMetrics treats a cash flow as a zero coupon bond and subjects it to two assumptions: (1) There is no expected change in the market value of such a bond, and (2) the volatility of the bond's market value scales up with the square root of the time horizon. In reality, the bond's market value systematically increases toward its par value (the "pull to par" effect), and its daily volatility decreases as it moves closer to par (the "roll down" effect). The two assumptions imply that the cash flow is treated as a generic bond (a bond whose maturity is always the same) rather than as an instrument whose maturity decreases with time. While this leads to an accurate depiction of the risk of the future cash flow for short forecast horizons, for longer horizons, it can result in a significant overstatement of risk.

Suppose that as of today, a USD based investor currently holds a one-year USD money market deposit and is interested in computing a Value-at-Risk estimate of this instrument over a 3-month forecast horizon. That is, the investor would like to know the maximum loss on this deposit (at a 95% confidence level) if he held the deposit for 3 months. To compute the risk of this position we compute the VaR of holding 9-month deposit with a forecast horizon 3-months. In other words, we are measuring the volatility on the 9-month deposit over a 3-month forecast horizon. To do this we use the current 9-month money market rates. This addresses the "roll down effect". In addition, the expected value of holding a one-year deposit for 3 months is not zero. Instead, the mean return is

[9] This section is based on the article by Christopher C. Finger, "Accounting for the "pull to par" and "roll down" for RiskMetrics cash flows", *RiskMetrics Monitor*, September 16, 1996.

non-zero reflecting the pull-to-par phenomenon. Chart 6.17 presents a visual description of the situation.

Chart 6.17
VaR horizon and maturity of money market deposit

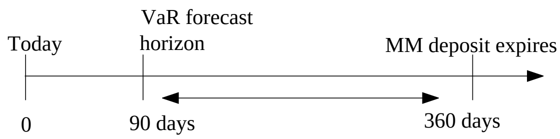

In general, the methodology to measure the VaR of a future cash flow(s) that occurs in T days over a forecast horizon of t days ($t < T$) is as follows.

1. Use the T-t rate, y_{T-t}, to discount the cash flow occurring in T days time. Denote the present value of this cash flow by V_{T-t}
2. Compute VaR as $V_{T-t} \cdot \left(\sigma_{T-t}\sqrt{t} \right)$.

Note that in the preceding example, $T = 360$, $t = 90$, y_{T-t} is the 270-day rate and σ_{T-t} is the standard deviation of the distribution of returns on the 270-day rate.

6.3.2.2 Equity positions
The market risk of the stock, VaR_t, is defined as the market value of the investment in that stock, V_t, multiplied by the price volatility estimate of that stock's returns, $1.65\sigma_t$.

[6.31] $$VaR_t = V_t \cdot 1.65\sigma_t$$

Since RiskMetrics does not publish volatility estimates for individual stocks, equity positions are mapped to their respective local indices. This methodology is based upon the principles of single-index models (the Capital Asset Pricing Model is one example) that relate the return of a stock to the return of a stock (market) index in order to attempt to forecast the correlation structure between securities. Let the return of a stock, r_t, be defined as

[6.32] $$r_t = \beta_t r_{m,t} + \alpha_t + \varepsilon_t$$

where

r_t = return of the stock
$r_{m,t}$ = the return of the market index
β_t = beta, a measure of the expected change in r_t given a change in $r_{m,t}$
α_t = the expected value of a stock's return that is firm-specific
ε_t = the random element of the firm-specific return with $E[\varepsilon_t] = 0$ and $E[\varepsilon_t]^2 = \sigma^2_{\varepsilon_t}$

As such, the returns of an asset are explained by market-specific components $(\beta_t R_{m,t})$[3] and by stock-specific components $(\alpha_t + \varepsilon_t)$. Similarly, the total variance of a stock is a function of the market- and firm-specific variances.

Since the firm-specific component can be diversified away by increasing the number of different equities that comprise a given portfolio, the market risk, VaR_t, of the stock can be expressed as a function of the stock index

[6.33] $$\sigma_t^2 = \beta_t^2 \sigma_{m,t}^2 + \sigma_{\varepsilon_t}^2 .$$

Substituting Eq. [6.33] into Eq. [6.31] yields

[6.34] $$VaR_t = V_t \cdot \beta_t \cdot 1.65\sigma_{m,t} ,$$

where

$$1.65\sigma_{m,t} = \text{the RiskMetrics VaR estimate for the appropriate stock index} .$$

As with individual stocks, Eq. [6.34] should also be used to calculate the VaR for positions that consist of issue that themselves are part of the EMBI+.

6.3.3 Delta-gamma VaR methodology (for portfolios containing options)

In this section we describe a methodology known as delta-gamma that allows users to compute the Value-at-Risk of a portfolio that contains options. Specifically, we provide a methodology to incorporate the delta, gamma and theta of individual options in the VaR calculation. We explain this methodology by first showing how it applies to a single option and then to a portfolio that contains three options. To keep our examples simple, we assume that each option is a function of one cash flow. In other words, we can write the return on each option as

[6.35] $$\hat{r}_{1,t} = \tilde{\delta}_1 r_{1,t} + 0.5\tilde{\Gamma}_1 r_{1,t}^2 + \tilde{\theta}_1 n$$

where
$$\begin{aligned} \tilde{\delta}_1 &= \eta_1 \delta_1 \\ \tilde{\Gamma}_1 &= \eta_1 P_{1,t} \Gamma_1 \\ \tilde{\theta}_1 &= \theta_1 / V_1 \\ n &= \text{VaR forecast horizon} \\ V_1 &= \text{option's premium} \end{aligned}$$

For a complete derivation of Eq. [6.35], see Appendix D. Similarly, we can write the returns on the other two options as

[6.36] $$\hat{r}_{2,t} = \tilde{\delta}_2 r_{2,t} + 0.5\tilde{\Gamma}_2 r_{2,t}^2 + \tilde{\theta}_2 n \text{ and } \hat{r}_{3,t} = \tilde{\delta}_3 r_{3,t} + 0.5\tilde{\Gamma}_3 r_{3,t}^2 + \tilde{\theta}_3 n$$

Let's begin by demonstrating the effect of incorporating gamma and theta components on the return distribution of the option. We do so by comparing the statistical features on the return on option 1, $\hat{r}_{1,t}$, and the return of its underlying cash flow, $r_{1,t}$. Recall that RiskMetrics assumes that the returns on the underlying assets, $r_{1,t}$, are normally distributed with mean 0 and variance

$\sigma^2_{1,t}$. Table 6.4 shows the first four moments[10]—the mean, variance, skewness, and kurtosis—for $\widehat{r}_{1,t}$ and $r_{1,t}$.

Table 6.4
Statistical features of an option and its underlying return

Statistical parameter	**Option**	**Underlying**
Return	$\widehat{r}_{1,t}$	$r_{1,t}$
Mean	$0.5\tilde{\Gamma}\sigma^2_{1,t} + \tilde{\theta}_1 n$	0
Variance	$\tilde{\delta}_1^2\sigma^2_{1,t} + 0.5\tilde{\Gamma}^2\sigma^4_{1,t}$	$\sigma^2_{1,t}$
Skewness	$3\tilde{\delta}_1^2\tilde{\Gamma}_1\sigma^4_{1,t} + \tilde{\Gamma}_1^3\sigma^6_{1,t}$	0
Kurtosis	$12\tilde{\delta}_1^2\tilde{\Gamma}_1^2\sigma^6_{1,t} + 3\tilde{\Gamma}_1^4\sigma^8_{1,t} + 3\sigma^4_{1,t}$	$3\sigma^4_{1,t}$

The results presented in Table 6.4 point to three interesting findings.

- First, even though it is assumed that the return on the underlying has a zero mean return, this is not true for the option's return unless both gamma and theta are zero. Also, the sign of the option's mean will be determined by the relative magnitudes and signs of both gamma and theta and whether one is long or short the option.

- Second, the variance of the return on the option differs from the variance of the return on the underlying instrument by the factor $\left(\tilde{\delta}_1^2 + 0.5\tilde{\Gamma}^2\sigma^2_{1,t}\right)$.

- And third, depending on whether one is long or short the option determines whether the return on the option distribution is negatively or positively skewed. To see this, on a short option position, $V_1 < 0$ and therefore $\tilde{\Gamma}_1 < 0$. Consequently, the term $\tilde{\Gamma}^3$ in the skewness expression will be negative. As an example of this point, Chart 6.18 shows the probability density functions for long and short options positions (along with the normal probability curve).

[10] See Section 4.5.2.1 for the definition of these moments.

Chart 6.18
Long and short option positions
negative and positive skewness

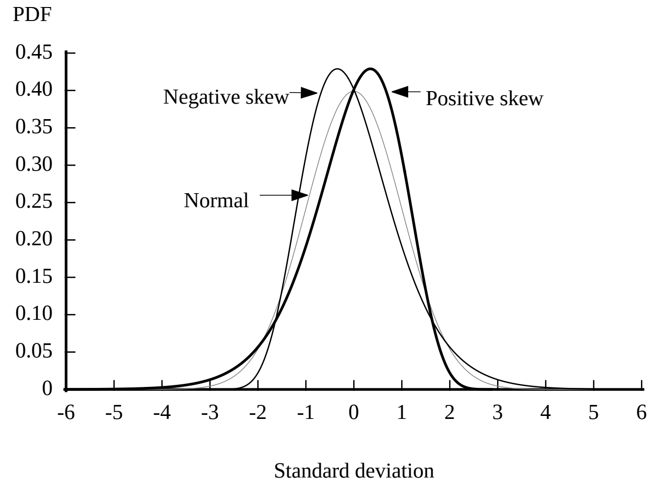

Note that in variance and kurtosis, the sign of $\tilde{\Gamma}$ is immaterial since in these expressions $\tilde{\Gamma}$ is raised to an even power.

Now, to determine the numerical values of the moments presented in Table 6.4 we need estimates of $\tilde{\delta}_1$, $\tilde{\Gamma}_1$, $\tilde{\theta}_1$ and $\sigma^2_{1,t}$. Estimates of the first three parameters are easily found by applying a Black-Scholes type valuation model. The variance, $\sigma^2_{1,t}$, is given in the RiskMetrics data sets.

Having obtained the first four moments of $\widehat{r}_{1,t}$'s distribution, we find a distribution that has the same moments but whose distribution we know exactly. In other words, we match the moments of $\widehat{r}_{1,t}$'s distribution to one of a set of possible distributions known as Johnson distributions. Here, "matching moments" simply means finding a distribution that has the same mean, standard deviation, skewness and kurtosis as $\widehat{r}_{1,t}$'s. The name Johnson comes from the statistician Norman Johnson who described a process of matching a particular distribution to a given set of moments.

Matching moments to a family of distributions requires that we specify a transformation from the option's return $\widehat{r}_{1,t}$ to a return, $r_{1,t}$, that has a standard normal distribution. For example, Johnson (1949) suggested the general transformation

[6.37] $$r_{1,t} = a + bf\left(\frac{\widehat{r}_{1,t} - c}{d}\right)$$

where $f(\ )$ is a monotonic function and a, b, c and d are parameters whose values are determined by $\widehat{r}_{1,t}$'s first four moments. In addition to the normal distribution, the Johnson family of distributions consists of three types of transformations.

[6.38] $$r_{1,t} = a + b\log\left(\frac{\widehat{r}_{1,t} - c}{d}\right) \quad \text{(Lognormal)}$$

[6.39] $$r_{1,t} = a + b\sinh^{-1}\left(\frac{\widehat{r}_{1,t} - c}{d}\right) \quad \text{(Unbounded)}$$

and

[6.40] $$r_{1,t} = a + b\log\left(\frac{\tilde{r}_{1,t} - c}{c + d - \tilde{r}_{1,t}}\right) \quad \text{(Bounded)}$$

with the restriction $(c < \tilde{r}_{1,t} < c + d)$.

To find estimates of a, b, c and d, we apply a modified version of Hill, Hill and Holder's (1976) algorithm.[11] Given these estimates we can calculate any percentile of $\tilde{r}_{1,t}$'s distribution given the corresponding standard normal percentile (e.g., –1.65). This approximate percentile is then used in the VaR calculation. For example, suppose that we have estimates of $\tilde{\delta}_1$, $\tilde{\Gamma}_1$, $\tilde{\theta}_1$ and $\sigma^2_{1,t}$ and that they result in the following moments: mean = 0.2, variance = 1, skewness coefficient = 0.75 and kurtosis coefficient = 7. Note that these numbers would be derived from the formulae presented in Table 6.3. Applying the Hill et. al. algorithm we find that the selected distribution is "Unbounded" with parameter estimates: a = –0.582, b = 1.768, c = –0.353, and d = 1.406.

Therefore, the percentile of $\tilde{r}_{1,t}$'s distribution is based on the transformation

[6.41] $$\tilde{r}_{1,t} = \sinh\left(\frac{(r_{1,t} - a)}{b}\right) \cdot d + c$$

Setting $r_{1,t} = -1.65$, the estimated 5th percentile of $\tilde{r}_{1,t}$'s distribution is −1.45. That is, the option's fifth percentile is increased by 0.20. In this hypothetical example, the incorporation of gamma and theta reduces the risk relative to holding the underlying.

We now show that it is straightforward to compute the VaR of a portfolio of options. In particular, we show this for the case of a portfolio that contains three options. We begin by writing the portfolio return as

[6.42] $$\tilde{r}_{p,t} = \omega_1 \tilde{r}_{1,t} + \omega_2 \tilde{r}_{2,t} + \omega_3 \tilde{r}_{3,t}$$

$$\text{where } \omega_i = \frac{V_i}{\sum_{i=1}^{3} V_i}$$

To compute the moments of $\tilde{r}_{p,t}$'s distribution we need the RiskMetrics covariance matrix, Σ, of the underlying returns $\{r_{1,t}, r_{2,t}, r_{3,t}\}$, and the delta, gamma and theta cash flow vectors that are defined as follows:

[6.43] $$\tilde{\delta} = \begin{bmatrix} \tilde{\delta}_1 \\ \tilde{\delta}_2 \\ \tilde{\delta}_3 \end{bmatrix}, \quad \tilde{\Gamma} = \begin{bmatrix} \tilde{\Gamma}_1 & 0 & 0 \\ 0 & \tilde{\Gamma}_2 & 0 \\ 0 & 0 & \tilde{\Gamma}_3 \end{bmatrix}, \text{ and } \tilde{\theta} = \begin{bmatrix} \tilde{\theta}_1 \\ \tilde{\theta}_2 \\ \tilde{\theta}_3 \end{bmatrix}$$

We find the 5th percentile $\tilde{r}_{p,t}$'s distribution the same way we found the 5th percentile of $\tilde{r}_{1,t}$'s distribution, as shown previously. The only difference is that now the expressions for the four moments are more complicated. For example, the mean and variance of the portfolio return are

[11] These original algorithms (numbers 99 and 100) are available in their entirety on the Web at the StatLib—Griffiths and Hill Archive. The URL is http://lib.stat.cmu.edu/griffiths-hill/.

[6.44] $$\mu_{p,t} = 0.5 \cdot \text{trace}\,[\tilde{\Gamma}\Sigma] + \sum_{i=1}^{3} \tilde{\theta}_i \text{ and}$$

[6.45] $$\sigma^2_{p,t} = \tilde{\delta}'\Sigma\tilde{\delta} + 0.5 \cdot \text{trace}\left[(\tilde{\Gamma}\Sigma)^2 \right]$$

where “trace” is an operator that sums the diagonal elements of a matrix.

The delta-gamma methodology described in this section extends to options that have more than one underlying cash flow (e.g., bond option). We have presented a simple example purposely to facilitate our exposition of the methodology. See, Appendix D for an assessment of the methodology.

Finally, the methodologies presented in Section 6.3 do not require simulation. All that is necessary for computing VaR is a covariance matrix, financial parameters (such as delta, gamma and theta) and position values. In the next chapter we present a methodology known as structured Monte Carlo that computes VaR by first simulating future paths of financial prices and/or rates.

6.4 Examples

In this section we present nine examples of VaR calculations for the various instruments discussed in this chapter. Note that the diskette symbol placed to the left of each example means that the example appears on the enclosed diskette at the back of the book.

Ex. 6.1 Government bond mapping of a single cash flow

Suppose that on March 27, 1995, an investor owns FRF 100,000 of the French OAT benchmark 7.5% maturing in April 2005. This bond pays coupon flows of FRF 7,500 each over the next 10 years and returns the principal investment at maturity. One of these flows occurs in 6.08 years, between the standard vertices of 5 and 7 years (for which volatilities and correlations are available).

All the data required to compute the cash flow map is readily available in the RiskMetrics data sets except for the price volatility ($1.65\sigma_{6.08}$) of the original cash flow's present value. This must be interpolated from the price volatilities already determined for the RiskMetrics vertices.

Applying the three conditions on page 118 and using Eqs. [6.10]–[6.14] with the RiskMetrics data in Table 6.5, we solve for the allocation α (and $(1-\alpha)$), and obtain the values $\alpha = 4.30$ and $\alpha = 0.4073$. Given the constraint that both of the allocated cash flows must have the same sign as the original cash flow, we reject the first solution, which would lead to a short position in the second proxy cash flow. As a result, our original cash flow of FRF 7,500, whose present value is FRF 4,774, must be mapped as a combination of a 5-year maturity cash flow of FRF 1,944 (40.73% of the original cash flow's PV) and a 7-year maturity cash flow of FRF 2,829 (59.27% of the original cash flow's PV).

The cash flow map is shown in Table 6.6.

Table 6.5
RiskMetrics data for 27, March 1995

RiskMetrics Vertex	Yield,%	P. Vol,* ($1.65\sigma_t$)	Yield Vol,† ($1.65\sigma_t$)	Correlation Matrix, ρ_{ij}	
				5yr	**7yr**
5yr	7.628	0.533	1.50	1.000	0.963
7yr	7.794	0.696	1.37	0.963	1.000

* P. Vol = price volatility, also called the VaR statistic.
† While this data is provided in the data set, it is not used in this calculation.

Table 6.6
RiskMetrics map of single cash flow

			Step 1*	Step 2*†		Step 3*		Step 4*	Step 5*
Coupon Flow	**Term**	**Yield,% (Actual)**	**Yield,% ($y_{6.08}$) (Interpolated)**	**$(PV)_{6.08}$‡**	**P. Vol, ($1.65\sigma_t$)§ (RiskMetrics)**	**P. Vol, ($1.65\sigma_{6.08}$)§ (Interpolated)**	**RiskMetrics Vertex**	**Allocation**	**RiskMetrics Cash flow**
		7.628			0.533		5yr	0.4073	1,950
7,500	6.08yr		7.717	4,774		0.624			
		7.794			0.696		7yr	0.3927	2,824

* Step from the mapping procedure on pages 119–121. Also, data in this column is calculated from the data in Table 6.5. Note that in Step 3 the price volatility, $1.65\sigma_{6.08}$, rather than the standard deviation alone, is computed.
† In this example $\hat{a} = 0.46$.
‡ PV = present value.
§ P. Vol = price volatility, also called the VaR statistic.

Ex. 6.2 Government bond mapping of multiple cash flows

A full set of positions can easily be mapped in the same fashion as the single cash flow in the last example.

The example below takes the instrument in Ex. 6.1, i.e., the 10-year French OAT benchmark on March 27, 1995, and decomposes all of the component cash flows according to the method described on pages 119–121, to create a detailed RiskMetrics cash flow map. Table 6.7 shows how the 100,000 French franc nominal position whose market value is FRF 97,400, is decomposed into nine representative present value cash flows. The table also shows the VaR for the cash flow at each RiskMetrics vertex and the diversified Value-at-Risk.

In this example, note that the first cash flow (on 25-Apr-95) occurs in less than one month's time relative to March 27, but is allocated at 100% weight to the 1-month RiskMetrics vertex. The reason for 100% allocation is that vertices shorter than one month are not defined in the RiskMetrics data sets.

Table 6.7
RiskMetrics map for multiple cash flows

Bond data	
Settlement	30-Mar
Principal	100,000
Price	97.4
Coupon	7.50
Basis	365

RiskMetrics™ vertices		**1m**	**1y**	**2y**	**3y**	**4y**	**5y**	**7y**	**10y**	**15y**
Yield volatility		7.00	3.16	2.10	1.74	1.63	1.50	1.37	1.36	1.29
Current yield		8.25	7.04	7.28	7.39	7.54	7.63	7.79	7.92	8.15
Price volatility		0.04	0.21	0.29	0.36	0.46	0.53	0.70	1.00	1.46
Correlation Matrix	1m	**1.00**	**0.75**	0.53	0.48	0.45	0.42	0.33	0.33	0.33
	1y	**0.75**	**1.00**	**0.88**	0.81	0.78	0.74	0.61	0.63	0.58
	2y	0.53	**0.88**	**1.00**	**0.99**	0.96	0.92	0.80	0.82	0.76
	3y	0.48	0.81	**0.99**	**1.00**	**0.98**	0.95	0.85	0.87	0.81
	4y	0.45	0.78	0.96	**0.98**	**1.00**	**0.99**	0.91	0.93	0.88
	5y	0.42	0.74	0.92	0.95	**0.99**	**1.00**	**0.96**	0.96	0.93
	7y	0.33	0.61	0.80	0.85	0.91	**0.96**	**1.00**	**1.00**	0.99
	10y	0.33	0.63	0.82	0.87	0.93	0.96	**1.00**	**1.00**	**0.99**
	15y	0.33	0.58	0.76	0.81	0.88	0.93	0.99	**0.99**	**1.00**

Date	Flow	Term	Yield	PV	Md. Dur	P.Vol	1m	1y	2y	3y	4y	5y	7y	10y	15y
25-Apr-95	7,500	0.071	8.204	7,456	0.066	0.032	7,456								
25-Apr-96	7,500	1.074	7.056	6,970	1.003	0.218		5,594	1,376						
25-Apr-97	7,500	2.074	7.284	6,482	1.933	0.292			5,780	703					
25-Apr-98	7,500	3.074	7.402	6,022	2.862	0.366				5,505	517				
25-Apr-99	7,500	4.074	7.543	5,577	3.788	0.463					5,105	472			
25-Apr-00	7,500	5.077	7.635	5,162	4.717	0.539						4,923	240		
25-Apr-01	7,500	6.077	7.720	4,773	5.641	0.624						1,944	2,829		
25-Apr-02	7,500	7.077	7.798	4,408	6.565	0.703							4,302	107	
25-Apr-03	7,500	8.077	7.855	4,072	7.488	0.805							2,589	1,483	
25-Apr-04	7,500	9.079	7.895	3,762	8.415	0.905							1,131	2,631	
25-Apr-05	107,500	10.079	7.919	49,863	9.340	1.004								49,019	844

RiskMetrics™ vertices	**1m**	**1y**	**2y**	**3y**	**4y**	**5y**	**7y**	**10y**	**15y**
Total Vertex Mapping	7,456	5,594	7,156	6,207	5,622	7,339	11,091	53,239	844
RiskMetrics™ Vertex VaR	3	12	20	22	26	39	77	530	12

Diversified Value at Risk	727	FRF over the next 24 hours
% of market value	0.7%	

Money market rate volatilities are used for vertices below 2 years. Government bond zero volatilities are used for 2-year and other vertices.

Ex. 6.3 Forward rate agreement cash flow mapping and VaR

A forward rate agreement is an interest-rate contract. It locks in an interest rate, either a borrowing rate (buying a FRA) or a lending rate (selling a FRA) for a specific period in the future. FRAs are similar to futures, but are over-the-counter instruments and can be customized for any maturity.

Because a FRA is a notional contract, there is no exchange of principal at the expiry date (i.e., the fixing date). If the rate is higher at settlement than the FRA rate agreed by the counterparties when they traded, then the seller of a FRA agrees to pay the buyer the present value of the interest rate differential applied to the nominal amount agreed upon at the time of the trade. The interest rate differential is between the FRA rate and the observed fixing rate for the period. In most cases this is the LIBOR rate for any given currency.

The general FRA pricing equation is given by

[6.46]
$$Y_{T_S \times T_L FRA} = \left[\frac{1 + \left(Y_L \cdot \frac{T_L}{B} \right)}{1 + \left(Y_S \cdot \frac{T_S}{B} \right)} - 1 \right] \cdot \frac{B}{T_L - T_S}$$

where

Y_L = yield of the longer maturity leg
T_L = maturity of the longer maturity leg in days
B = Basis (360 or 365)
Y_S = yield of the shorter maturity leg
T_S = maturity of the shorter maturity leg in days

In effect, FRAs allow market participants to lock in a forward rate that equals the implied break-even rate between money market term deposits.

Given that a FRA is a linear combination of money market rates, it is simple to express its degree of risk as a function of the combination of these rates.

Suppose that on January 6, 1995 you sold a 6x12 FRA on a notional 1 million French francs at 7.24%. This is equivalent to locking in an investment rate for 6 months starting in 6 months' time. The rate of 7.24% is calculated by combining the 6- and 12-month money market rates using the general pricing equation, Eq. [6.46], which can be rewritten as follows to reflect the no-arbitrage condition:

[6.47]
$$\left(1 + Y_{12m} \cdot \frac{365}{365} \right) = \left(1 + Y_{6m} \cdot \frac{181}{365} \right) \cdot \left(1 + Y_{6 \times 12 FRA} \cdot \frac{184}{365} \right)$$

where

Y_{6m}, Y_{12m} = 6- and 12-month French franc yields, respectively
Y_{6m}, Y_{12m} = 6 x 12 FRA rate

This FRA transaction is equivalent to borrowing FRF 1 million for 6 months on a discount basis (i.e., total liability of FRF 1 million in 6 months' time) and investing the proceeds (FRF 969,121) for 12 months. This combination can be mapped easily into the RiskMetrics vertices as shown in Table 6.8. The current present value of these two positions is shown in column (6). The Value-at-

Example 6.3 (continued)

Risk of each leg of the FRA over a 1-month horizon period, shown in column (9), is obtained by multiplying the absolute present value of the position by the monthly price volatility of the equivalent maturity deposit. The portfolio VaR is obtained by applying the 6- to 12-month correlation to the VaR estimate from column (9).

Table 6.8
Mapping a 6x12 short FRF FRA at inception

Observed data		**RiskMetrics data set**					**Calculated values**			
(1)	(2)	(3)	(4)		(5)		(6)	(7)	(8)	(9)
			Volatilities		Correlation matrix					
Cash flow	Term (mths.)	Yield,%	Yield	Price	6m	12m	Present value	RiskMetrics vertex	RiskMetrics cash flow	VaR estimate
−1,000,000	6	6.39	6.94	0.21	1.00	0.70	-969,121	6m	−969,121	2,081
1,036,317	12	6.93	7.42	0.48	0.70	1.00	969,121	12m	969,121	4,662
						Total	0		0	
						Portfolio VaR				3,530

One month into the trade, the mapping becomes somewhat more complex as the cash flows have now shorter maturities (the instrument is now in fact a 5x11 FRA). The 5-month cash flow must be mapped as a combination of 3-month and 6-month RiskMetrics vertices (Table 6.9), while the 11-month cash flow must be split between the 6-month and 12-month vertices.

Table 6.9
Mapping a 6x12 short FRF FRA held for one month

Observed data		**RiskMetrics data set**					**Calculated values**				
		Volatilities		Correlation matrix							
Cash flow	Term (mths.)	Yield	Price	3m	6m	12m	Yield,%	Present value	RiskMetrics vertex	RiskMetrics cash flow	VaR estimate
		6.77	0.1	1	0.81	0.67			3m	−302,232	−296
−1,000,000	5						6.12	−975,302			
		7.91	0.19	0.81	1.00	0.68			6m	−549,300	−1,048
1,036,317	11						6.65	976,894			
		7.14	0.41	0.67	0.68	1.00			12m	853,124	3,533
							Total	1,592		1,592	
							Portfolio VaR				2,777

Example 6.3 (continued)

One month into the trade, the change in market value of the contract is a positive FRF 1,592. This is well within the range of possible gains or losses predicted (with a 95% probability) by the previous month's Value-at-Risk estimate of FRF 3,530.

Unwinding a FRA, i.e., hedging out the interest rate risk between the FRA rate and the market rate, before maturity requires entering into a contract of opposite sign at dates that no longer qualify as standard maturities. If you wanted to unwind the position in this example one month after the dealing date, you would have to ask a quote to buy a 5x11 FRA, a broken dated instrument that is less liquid and therefore is quoted at higher bid-offer spreads. The rates in column (1) above do not take this into account. They were derived by interpolating rates between standard maturities. Actual market quotes would have been slightly less favorable, reducing the profit on the transaction. This risk is liquidity related and is not identified in the VaR calculations.

Ex. 6.4 Structured note

The basic concepts used to estimate the market risk of simple derivatives can be extended to more complex instruments. Suppose that in early 1994, when the market consensus was that German rates were to continue to decline, you had purchased a “one year index note linked to two year rates”. This 1-year instrument leveraged a view that the DEM 2-year swap rate in 1-year’s time would be below the forward rate measured at the time the transaction was entered into.

The characteristics of the instrument are shown in Table 6.10.

Table 6.10
Structured note specification

Issuer	Company A
Format	Euro Medium Term Note
Issue date	9 February 94
Maturity date	9 February 95
Issue price	100%
Amount	DEM 35,000,000
Coupon	5.10%
Strike	5.10%
Redemption value	100%+20*(Strike- 2-year DEM swap rate)

* Although these details are hypothetical, similar products were marketed in 1994.

While seemingly complex, this transaction is in fact little more (disregarding minor convexity issues) than a bond to which a leveraged long-dated FRA had been attached. As a holder of the note, you were long the 3-year swap rate and short the 1-year rate, with significant leverage attached to the difference.

Table 6.11 shows how the leveraged note can be decomposed into the cash flows of the two underlying building blocks:

- The 1-year DEM 35 million bond with a 5.10% coupon.

- The forward swap (2-year swap starting in one year). The forward principal cash flows of the swap are equal to 20 times the notional amount of the note divided by the PVBP (price value of a basis point) of a 2-year instrument, or 1.86 in this case. The forward coupons are equal to the forward principal times the coupon rate of 5.10%.

Table 6.11
Actual cash flows of a structured note

Term (years)	Bond		Swap		Total cash flow
	Principal	Coupon	Principal	Coupon	
1	35,000,000	1,785,000	–376,996,928		–340,211,928
2				19,226,843	19,226,843
3			376,996,928	19,226,843	396,223,771

Combining the bond and the swap creates three annual cash flows where the investor is short DEM 340 million in the 1-year, and long DEM 19 and DEM 396 million in the 2- and 3-year maturities. At issue, the market value of these cash flows is equal to DEM 35 million, the instrument’s issue price.

Example 6.4 (continued)

Each of the three cash flows is mapped to RiskMetrics vertices to produce the cash flow map shown in Table 6.12.

Table 6.12
VaR calculation of structured note
One month forecast horizon

Observed data		**RiskMetrics data set**					**Calculated values**			
				Correlation matrix						
Cash flow	Term (years)	Yield,%	Price volatility	1y	2y	3y	Present value	RiskMetrics vertex	RiskMetrics cash flow	VaR estimate
–340,211,928	1	5.48	0.33	1.00	0.46	0.43	–322,536,906	1y	322,536,906	1,067,597
19,226,843	2	5.15	0.46	0.46	1.00	0.95	17,389,594	2y	17,389,594	79,644
396,223,772	3	5.22	0.68	0.43	0.95	1.00	340,147,312	3y	340,147,312	2,309,600
						Total	35,000,000			
						Portfolio VaR				2,155,108

Using the appropriate volatilities and correlations as of February 9, 1994, the Value-at-Risk of such a position over a 1-month horizon was around DEM 2.15 million.

One month into the life of the instrument on March 9, the mapping and risk estimation could have been repeated using updated interest rates as well as updated RiskMetrics volatilities and correlations. Table 6.13 shows the result.

Table 6.13
VaR calculation on structured note
One-month into life of instrument

Observed data		**RiskMetrics data set**						**Calculated values**			
				Correlation matrix							
Cash flow	Term (years)	Yield,%	Price volatility	6m	1y	2y	3y	Present value	RiskMetrics vertex	RiskMetrics cash flow	VaR estimate
			0.16	1.00	0.83	0.58	0.54		6m	–43,218,017	68,023
–340,211,928	0.9	5.53						–975,302			
			0.35	0.83	1.00	0.58	0.54		1y	–277,979,823	960,250
19,226,843	1.9	5.53						17,399,085			
			0.65	0.58	0.58	0.94	1.00		2y	39,059,679	252180
19,226,843	2.9	5.68						337,230,727			
			1.03	0.54	0.54	0.94	1.00		3y	312,934,965	3,218,971
							Total	30,976,801		–4,203,199	
							Portfolio VaR				3,018,143

The movement in market rates led the market value of the note to fall by over DEM 4 million, twice the maximum amount expected to happen with a 95% probability using the previous month's RiskMetrics volatilities and correlations. Why?

Example 6.4 (continued)

Chart 6.19 shows the 3-year DEM swap rate moved 45+bp from 5.22% to slightly above 5.68% during the month—twice the maximum amount expected with a 95% probability (4.56% × 5.22%; i.e., 23 basis points). This was clearly a large rate move. The RiskMetrics volatility estimate increased from 4.56% to 6.37% as of March 9. This reflects the rapid adjustment to recent observations resulting from the use of an exponential moving average estimation method. Correspondingly the VaR of the structured note increased 44% over the period to just over DEM 3 million.

Chart 6.19
DEM 3-year swaps in Q1-94

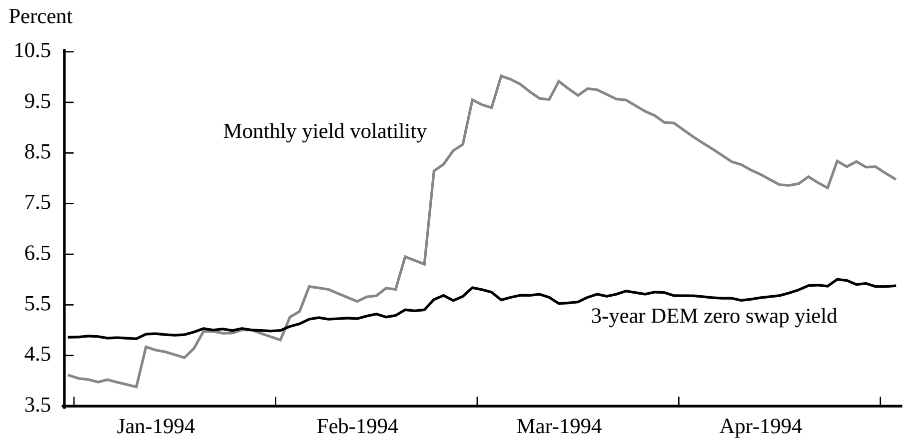

The message in these examples is that with proper cash flow mapping, the risks in complex derivatives can be easily estimated using the RiskMetrics methodology and data sets.

Ex. 6.5 Interest-rate swap

Investors enter into swaps to change their exposure to interest rate uncertainty by exchanging interest flows. In order to understand how to map its cash flows, a swap should be thought of as a portfolio of one fixed- and one floating-rate transaction. Specifically, the fixed leg of a swap is mapped as if it were a bond, while the floating leg is considered to be a FRN.

Market risk estimation is straightforward if the value of each leg is considered separately. The fixed leg exposes an investor to interest rate variability as would a bond. Since the floating leg is valued as if it were a FRN, if interest rates change, then forward rates used to value the leg change and it will revalue to par. Once a floating payment is set, the remaining portion of the floating leg will revalue to par, and we need only consider interest rate exposure with respect to that set cash flow. The details of this will be provided in a forthcoming edition of the *RiskMetrics Monitor.*

Consider the following example. A company that enters into a 5-year USD interest-rate swap pays 9.379% fixed and receives floating cash flows indexed off of 1-year USD LIBOR flat on a notional amount of USD 1,000,000. For simplicity, the reset/payment dates are annual. Table 6.14 presents the data used to estimate the market risk of this transaction.

Table 6.14
Cash flow mapping and VaR of interest-rate swap

Observed data				**RiskMetrics data set**						**Calculated values**
		Cash flow			Correlation matrix					
Term	Zero rate	Fixed	Floating	Price volatility, %	1yr	2yr	3yr	4yr	5yr	Net present value, USD
1yr	8.75	−86,247	1,000,000	0.027	1.000	0.949	0.933	0.923	0.911	913,753
2yr	9.08	−85,986		0.067	0.949	1.000	0.982	0.978	0.964	−85,986
3yr	9.24	−85,860		0.112	0.933	0.982	1.0000	0.995	0.984	85,860
4yr	9.34	−85,782		0.149	0.923	0.978	0.995	1.000	0.986	85,782
5yr	9.42	−999,629		0.190	0.911	0.964	0.984	0.986	1.000	−999,629
							Total			−343,505
							Portfolio VaR			1,958

Ex. 6.6 Foreign exchange forward

Below is an example of how to calculate the market risk of buying a 1-year 153,000,000 DEM/USD foreign exchange forward. Note that buying a DEM/USD foreign exchange forward is equivalent to borrowing US dollars for 1-year (short money market position) and using them to purchase Deutsche marks in one year's time (short foreign exchange position). We take the holding period to be one day. Based on a 1-day volatility forecast, the foreign exchange risk, in USD, is \$904,922 (94,004,163 × 0.963%) as shown in Table 6.15. The interest rate risk is calculated by multiplying the current market value of each 12-month leg (the short in USD and the long in DEM) times its respective interest rate volatility. Therefore, the Value-at-Risk for a 1-day holding period is \$912,880.

Table 6.15
VaR on foreign exchange forward

Observed data			**RiskMetrics data set**					**Calculated values**	
					Correlation matrix				
Instrument	Cash flow	Term (years)	Yield,%	Price volatility	DEM FX	DEM 1y	USD 1y	Present value, USD	VaR estimate
DEM Spot FX	—			0.963	1.0000	–0.0035	–0.0042	–94,004,163	–904,922
DEM 1y	153,000,000	1	6.12	0.074	–0.0035	1.0000	0.1240	94,004,163	45,855
USD 1y	–99,820,670	1	6.65	0.116	–0.0042	0.1240	1.0000	–94,004,163	–108,624
							Total	94,004,163	
							Portfolio VaR		3,530

Ex. 6.7 Equity

Consider a three-asset portfolio in which an investor holds stocks ABC and XYZ (both U.S. companies) as well as a basket of stocks that replicate the S & P 500 index. The market risk of this portfolio, VaR_p, is

$$VaR_p = (V_{ABC} \cdot 1.65\sigma_{R_{ABC}}) + (V_{XYZ} \cdot 1.65\sigma_{R_{XYZ}}) + V_{SP500} \cdot 1.65\sigma_{R_{SP500}} \quad [6.48]$$

Rewriting this equation in terms of Eq. [6.48]

$$VaR_s = V_s \cdot \beta_s \cdot 1.65\sigma_{R_M} \quad [6.49]$$

where

$1.65\sigma_{R_M}$ = the RiskMetrics volatility estimate for the appropriate stock index,

yields

$$VaR_p = (V_{ABC} \cdot \beta_{ABC} \cdot 1.65\sigma_{R_{SP500}}) + (V_{XYZ} \cdot \beta_{XYZ} \cdot 1.65\sigma_{R_{SP500}}) + (V_{SP500} \cdot 1.65\sigma_{R_{SP500}}) \quad [6.50]$$

Factoring the common term and solving for the portfolio VaR results in

$$\begin{aligned} VaR_p &= 1.65\sigma_{R_{SP500}}[(V_{ABC} \cdot \beta_{ABC}) + (V_{XYZ} \cdot \beta_{XYZ}) + V_{SP500}] \\ &= 4.832\%\,[1,000,000\,(0.5) + (1,000,000)\,(1.5) + 1,000,000] \\ &= 4.832\%\,(3,000,000) \\ &= \text{USD } 144,960 \end{aligned} \quad [6.51]$$

The methodology for estimating the market risk of a multi-index portfolio is similar to the process above and takes into account correlation among indices as well as foreign exchange rates. Since all positions must be described in a base or “home” currency, you need to account for foreign exchange risk.

Ex. 6.8 Commodity futures contract

Suppose on July 1, 1994 you bought a 6-month WTI future on a notional USD 18.3 million (1 million barrels multiplied by a price of USD 18.30 per barrel). The market and RiskMetrics data for that date are presented in Table 6.16.

Table 6.16
Market data and RiskMetrics estimates as of trade date July 1, 1994

			WTI future		Correlation matrix	
Vertex	LIBOR	Term	Price	Volatility	3m	6m
3m	5.563	0.250		10.25	1.000	
6m	5.813	0.500	18.30	9.47	0.992	1.000

The initial Value-at-Risk for a 1-month horizon is approximately USD 1.7 million. This represents the maximum amount, with 95% confidence, that one can expect to lose from this transaction over the next 25 business days. Since the flow occurs in 6 months, the entire position is mapped to the 6-month WTI vertex, therefore calculating the Value-at-Risk of this transaction on the trade date is simply

$$VaR_{6\text{m Future}} = \textit{PV of cash flow} \cdot \textit{RiskMetrics volatility estimate}$$

[6.52]
$$= \left[\frac{18,300,000}{1 + \left(\frac{5.813\%}{100}\right) \cdot 0.5} \right] \cdot 9.47\%$$

$$= \text{USD 1.7 million}$$

One month into the trade, the cash flow mapping becomes slightly more complex. Table 6.17 shows the new VaR of this transaction.

Table 6.17
Cash flow mapping and VaR of commodity futures contract

Term	Zero rate	Cash flow (PV)	Price volatility, %	Correlation matrix 3m	Correlation matrix 6m	RiskMetrics vertex	RiskMetrics cash flow
	4.810		6.212	1.00	0.992	3m	13,084,859
4m	5.068	17,924,465					
	5.190		5.739	0.992	1.00	6m	4,839,605
					Portfolio VaR		1,417,343

Ex. 6.9 Delta-gamma methodology

Consider the situation where a USD (US dollar) based investor currently holds a USD 1million equivalent position in a French government bond that matures in 6 years and a call option on Deutsche marks that expires in 3 months. Since market risk is measured in terms of a portfolio's return distribution, the first step to computing VaR is to write down an expression for this portfolio's return, R_P, which consists of one French government bond and one foreign exchange option. Here, return is defined as the one-day relative price change in the portfolio's value. The return on the portfolio is given by the expression

$$\widehat{r}_p = \widehat{r}_o + r_B \quad [6.53]$$

where r_o is the return on the option, and r_B is the return on the French government bond.

We now provide a more detailed expression for the returns on the bond and option. Since the cash flow generated by the bond does not coincide with a specific RiskMetrics vertex, we must map it to the two nearest RiskMetrics vertices. Suppose we map 49% of the cash flow that arrives in 6 years' time to the 5-year vertex and 51 percent of the cash flow to the 7-year vertex. If we denote the returns on the 5 and 7-year bonds by r_5 and r_7, respectively, we can write the return on the French government bond as

$$r_B = 0.49r_5 + 0.51r_7 \quad [6.54]$$

Writing the return on the option is more involved. We write the return on the option as a function of its delta, gamma and theta components. The one-day return on the option is given by the expression[12]

$$\widehat{r}_0 = \alpha\delta r_{\mathrm{USD/DEM}} + 0.5\alpha\Gamma P_{\mathrm{USD/DEM}} r^2_{\mathrm{USD/DEM}} + V^{-1}\theta n \quad [6.55]$$

where

$r_{\mathrm{USD/DEM}}$ is the one-day return on the DEM/USD exchange rate

$P_{\mathrm{USD/DEM}}$ is the spot position in USD/DEM

V is the price of the option, or premium.

α is the ratio of $P_{\mathrm{USD/DEM}}$ to V. The parameter α measures the leverage from holding the option.

δ is the "delta" of the option. Delta measures the change in the value of the option given a change in the underlying exchange rate.

Γ is the "gamma" of the option. Gamma measures the change in δ given a change in the underlying exchange rate.

θ is the "theta" of the option. Theta measures the change in the value of the option for a given change in the option's time to expiry.

n is the forecast horizon over which VaR is measured. In this example *n* is 1 for one day.

[12] We derive this expression in Appendix D.

Example 6.9 (continued)

We can now write the return on the portfolio as

$$[6.56] \qquad \hat{r}_p = 0.49 r_5 + 0.51 R_7 + \alpha\delta r_{USD/DEM} + 0.5\alpha\Gamma P_{USD/DEM} r^2_{USD/DEM} + V^{-1}\theta n$$

In particular, we find the first four moments of r_p 's distribution that correspond to the mean, variance, skewness and kurtosis (a measure of tail thickness). These moments depend only on the price of the option, the current market prices of the underlying securities, the option's "greeks" δ, Γ, θ, and the RiskMetrics covariance matrix, Σ. In this example, Σ is the covariance matrix of returns r_5, r_7 and $r_{USD/DEM}$.

Let's take a simple hypothetical example to describe the delta-gamma methodology. Table 6.18 presents the necessary statistics on the bond and option positions to apply delta-gamma.

Table 6.18
Portfolio specification

Bond	**Option**
$\sigma_5 = 0.95\%$	$\sigma_{USD/DEM} = 1\%$
$\sigma_7 = 1\%$	$\delta = 0.9032$
$\rho_{5,7} = 0.85$	$\Gamma = 0.0566$
$P_B = \text{USD } 100$	$\theta = -0.9156$
	V = USD 3.7191
	$P_{USd/DEM} = \text{USD } 346.3$
Portfolio P*V* = USD 103.719	

To compute VaR we require the covariance matrix

$$[6.57] \qquad \Sigma = \begin{bmatrix} \sigma^2_{6y} & \sigma^2_{6y,\,USD/DEM} \\ \sigma^2_{6y,\,USD/DEM} & \sigma^2_{USD/DEM} \end{bmatrix}$$

which, when using the information in Table 6.18 yields

$$[6.58] \qquad \Sigma = \begin{bmatrix} 0.009025 & 0.008075 \\ 0.008075 & 0.01000 \end{bmatrix} \text{ (in percent)}$$

Also, we need the cash flows corresponding to the delta components of the portfolio,

$$[6.59] \qquad \tilde{\delta} = \begin{bmatrix} 100 \\ 81.352 \end{bmatrix} \begin{matrix} \text{Delta cash flow on bond} \\ \text{Delta cash flow on FX option} \end{matrix}$$

the cash flows corresponding to the gamma components of the portfolio,

Example 6.9 (continued)

[6.60] $$\tilde{\Gamma} = \begin{bmatrix} 0 \\ 1708.47 \end{bmatrix} \begin{matrix} \text{Zero gamma cash flow on bond} \\ \text{Gamma cash flow for FX option} \end{matrix}$$

and, the cash flows corresponding to the theta components of the portfolio.

[6.61] $$\tilde{\theta} = \begin{bmatrix} 0 \\ -0.2462 \end{bmatrix} \begin{matrix} \text{Zero theta cashflow on bond} \\ \text{Theta cashflow for FX option} \end{matrix}$$

The implied moments of this portfolio's distribution are presented in Table 6.19.

Table 6.19
Portfolio statistics

Moments	Including theta	Excluding theta
mean	–0.1608	0.0854
variance	2.8927	2.8927
skewness coefficient	0.2747	0.2747
kurtosis coefficient	3.0997	3.0997

Based on the information presented above, VaR estimates of this portfolio over a one-day forecast horizon are presented in Table 6.20 for three confidence levels. For comparison we also present VaR based on the normal model and VaR that excludes the theta effect.

Table 6.20
Value-at-Risk estimates (USD)
one-day forecast horizon: total portfolio value is 103.719

VaR percentile	Normal	Delta-gamma (excluding theta)	Delta-gamma (including theta)
5.0%	–2.799	–2.579	–2.826
2.5%	–3.325	–3.018	–3.265
1.0%	–3.953	–3.523	–3.953

Chapter 7. Monte Carlo

Chapter 7. Monte Carlo

Christopher C. Finger
Morgan Guaranty Trust Company
Risk Management Research
(1-212) 648-4657
finger_christopher@jpmorgan.com

In the previous chapter, we illustrated how to combine cash flows, volatilities, and correlations analytically to compute the Value-at-Risk for a portfolio. We have seen that this methodology is applicable to linear instruments, as well as to non-linear instruments whose values can be well approximated by a Taylor series expansion (that is, by its "greeks").

In this chapter, we outline a Monte Carlo framework under which it is possible to compute VaR for portfolios whose instruments may not be amenable to the analytic treatment. We will see that this methodology produces an estimate for the entire probability distribution of portfolio values, and not just one risk measure.

The Monte Carlo methodology consists of three major steps:

1. **Scenario generation**—Using the volatility and correlation estimates for the underlying assets in our portfolio, we produce a large number of future price scenarios in accordance with the lognormal models described previously.
2. **Portfolio valuation—**For each scenario, we compute a portfolio value.
3. **Summary—**We report the results of the simulation, either as a portfolio distribution or as a particular risk measure.

We devote one section of this chapter to each of the three steps above.

To better demonstrate the methodology, we will consider throughout this section a portfolio comprised of two assets: a future cash flow stream of DEM 1M to be received in one year's time and an at the money put option with contract size of DEM 1M and expiration date one month in the future. Assume the implied volatility at which the option is priced is 14%. We see that our portfolio value is dependent on the USD/DEM exchange rate and the one year DEM bond price. (Technically, the value of the option also changes with USD interest rates and the implied volatility, but we will not consider these effects.) Our risk horizon for the example will be five days.

7.1 Scenario generation

We first recall the lognormal model which we assume for all underlying instruments. Consider a forecast horizon of t days. If an instrument's price today is P_0, and our estimate for the one day volatility of this instrument is σ, then we model the price of the instrument in t days by

$$P_t = P_0 e^{\sigma\sqrt{t}Y} \qquad [7.1]$$

where Y is a standard normal random variable. Thus, the procedure to generate scenarios is to generate standard normal variates and use Eq. [7.1] to produce future prices. The procedure for the multivariate case is similar, with the added complication that the Y's corresponding to each instrument must be correlated according to our correlation estimates.

In practice, it is straightforward to generate independent normal variates; generating arbitrarily correlated variates is more involved, however. Suppose we wish to generate n normal variates with unit variance and correlations given by the $n \times n$ matrix Λ. The general idea is to generate n independent variates, and then combine these variates is such a way to achieve the desired correlations. To be more precise, the procedure is as follows:

- Decompose Λ using the Cholesky factorization, yielding a lower triangular matrix A such that $\Lambda = AA'$. We provide details on this factorization below and in Appendix E.

- Generate an $n \times 1$ vector Z of independent standard normal variates.

- Let $Z = AY$. The elements of Z will each have unit variance and will be correlated according to Λ.

To illustrate the intuition behind using the Cholesky decomposition, consider the case where we wish to generate two variates with correlation matrix

[7.2] $$\Lambda = \begin{bmatrix} 1 & \rho \\ \rho & 1 \end{bmatrix}$$

The Cholesky factorization of Λ is given by

[7.3] $$A = \begin{bmatrix} 1 & 0 \\ \rho & \sqrt{1-\rho^2} \end{bmatrix}$$

(It is easy to check that $AA' = \Lambda$.) Now say that Y is a 2×1 vector containing independent standard normal random variables Y_1 and Y_2. If we let $Z = AY$, then the elements of Z are given by

[7.4a] $$Z_1 = Y_1 \text{ and}$$

[7.4b] $$Z_2 = \rho Y_1 + \sqrt{1-\rho^2}\,Y_2$$

Clearly, Z_1 has unit variance, and since Y_1 and Y_2 are independent, the variance of Z_2 is given by

[7.5] $$\rho^2 Var(Y_1) + \left(\sqrt{1-\rho^2}\right)^2 Var(Y_2) = 1$$

Again using the fact that Y_1 and Y_2 are independent, we see that the expected value of $Z_1 Z_2$ is just ρ, and so the correlation is as desired.

Note that it is not necessary to use the Cholesky factorization, since any matrix A which satisfies $\Lambda = AA'$ will serve in the procedure above. A singular value or eigenvalue decomposition would yield the same results. The Cholesky approach is advantageous since the lower triangular structure of A means that fewer operations are necessary in the AZ multiplication. Further, there exist recursive algorithms to compute the Cholesky factorization; we provide details on this in Appendix E. On the other hand, the Cholesky decomposition requires a positive-definite correlation matrix; large matrices obtained from the RiskMetrics data do not always have this property.

Using the procedure above to generate random variates with arbitrary correlations, we may generate scenarios of asset prices. For example, suppose we wish to model the prices of two assets t days into the future. Let $P_0^{(1)}$ and $P_0^{(2)}$ indicate the prices of the assets today, let σ_1 and σ_2 indicate the daily volatilities of the assets, and let ρ indicate the correlation between the two assets. To generate a future price scenario, we generate correlated standard normal variates Z_1 and Z_2 as outlined above and compute the future prices by

[7.6a] $$P_t^{(1)} = P_0^{(1)} e^{\sigma_1 \sqrt{t} Z_1} \text{ and}$$

[7.6b] $$P_t^{(2)} = P_0^{(2)} e^{\sigma_2 \sqrt{t} Z_2}$$

To generate a collection of scenarios, we simply repeat this procedure the required number of times.

For our example portfolio, the two underlying assets to be simulated are the USD/DEM exchange rate and the one year DEM bond price. Suppose that the current one year German interest rate is 10% (meaning the present value of a one year 1M DEM notional bond is DEM 909,091) and that the current USD/DEM exchange rate is 0.65. We take as the daily volatilities of these two assets $\sigma_{FX} = 0.0042$ and $\sigma_B = 0.0008$ and as the correlation between the two $\rho = -0.17$.

To generate one thousand scenarios for values of the two underlying assets in five days, we first use the approach above to generate one thousand pairs of standard normal variates whose correlation is ρ . Label each pair Z_{FX} and Z_B . We present histograms for the results in Chart 7.1. Note that the distributions are essentially the same.

Chart 7.1
Frequency distributions for Z_{FX} and Z_B
1000 trials

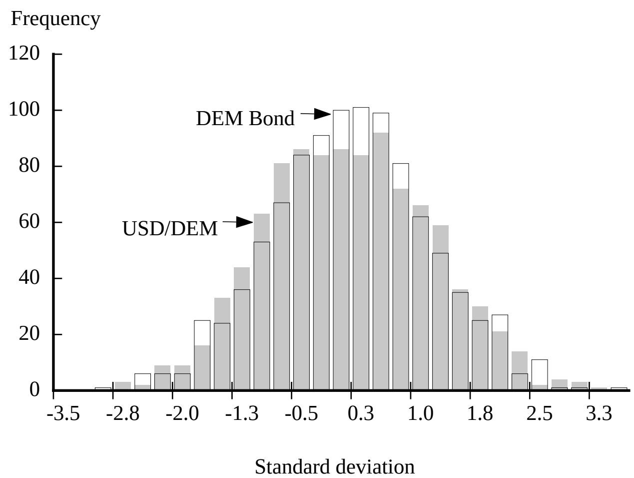

The next step is to apply Eq. [7.6a] and Eq. [7.6b]. This will create the actual scenarios for our assets. Thus, for each pair Z_{FX} and Z_B , we create future prices P_{FX} and P_B by applying

[7.7a] $$P_{FX} = 0.65 e^{0.0042 \times \sqrt{5} \times Z_{FX}}$$

and

[7.7b] $$P_B = 909,091 e^{0.0008 \times \sqrt{5} \times Z_B}$$

Of course, to express the bond price in USD (accounting for both the exchange rate and interest rate risk for the bond), it is necessary to multiply the bond price by the exchange rate in each scenario. Charts 7.2 and 7.3 show the distributions of future prices, P_B and P_{FX} , respectively, obtained by one thousand simulations. Note that the distributions are no longer bell shaped, and for the bond price, the distribution shows a marked asymmetry. This is due to the transformation we make from normal to lognormal variates by applying Eq. [7.7a] and Eq. [7.7b].

Chart 7.2
Frequency distribution for DEM bond price
1000 trials

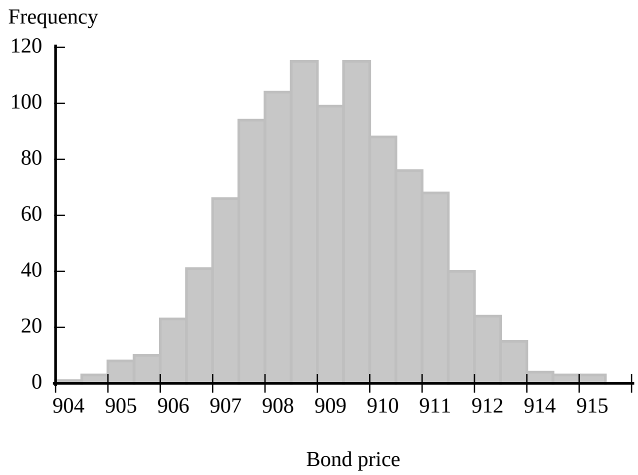

Chart 7.3
Frequency distribution for USD/DEM exchange rate
1000 trials

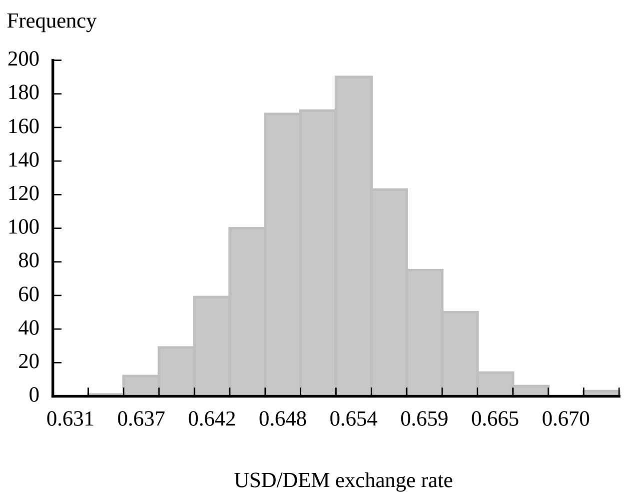

In Table 7.1, we present the first ten scenarios that we generate.

Table 7.1
Monte Carlo scenarios
1000 trials

USD/DEM	PV of cash flow (DEM)	PV of cash flow (USD)
0.6500	906,663	589,350
0.6540	907,898	593,742
0.6606	911,214	601,935
0.6513	908,004	591,399
0.6707	910,074	610,430
0.6444	908,478	585,460
0.6569	908,860	597,053
0.6559	906,797	594,789
0.6530	906,931	592,267
0.6625	920,768	603,348

Portfolio valuation

In the previous section, we illustrated how to generate scenarios of future prices for the underlying instruments in a portfolio. Here, we take up the next step, how to value the portfolio for each of these scenarios. We will examine three alternatives: full valuation, linear approximation, and higher order approximation. Each of the alternatives is parametric, that is, an approach in which the value of all securities in the portfolio is obtained through the values of its underlying assets, and differ only in their methods for valuing non-linear instruments given underlying prices.

Recall that at the current time, the present value of our cash flow is DEM 909,091, or USD 590,909. The value of the option is USD 10,479.

7.1.1 Full valuation

This is the most straightforward and most accurate alternative, but also the most intense computationally. We assume some type of pricing formula, in our case the Black-Scholes option pricing formula, with which we may value our option in each of the scenarios which we have generated. Say $V(S, K, \tau)$ gives the premium (in USD) associated with the option of selling one DEM given spot USD/DEM rate of S, strike rate of K, and expiration date τ years into the future. (Again, this function will also depend on interest rates and the implied volatility, but we will not model changes in these variables, and so will suppress them in the notation.)

In our example, for a scenario in which the USD/DEM rate has moved to R after five days, our option's value (in USD) moves from $1,000,000 \times V\left(0.65, 0.65, \frac{1}{12}\right)$ to $1,000,000 \times V\left(R, 0.65, \frac{1}{12} - \frac{5}{365}\right)$. The results of applying this method to our scenarios are displayed in Table 7.2. Note that scenarios with moderate changes in the bond price can display significant changes in the value of the option.

Table 7.2
Monte Carlo scenarios—valuation of option
1000 trials

	Value of option (USD)				
USD/DEM	PV of cash flow (USD)	Full	Delta	Delta/Gamma	Delta/Gamma/ Theta
0.6500	589,350	9,558	10,458	10,458	9,597
0.6540	593,742	7,752	8,524	8,644	7,783
0.6606	601,935	5,273	5,272	6,122	5,261
0.6513	591,399	8,945	9,831	9,844	8,893
0.6707	610,430	2,680	273	3,541	2,680
0.6444	585,460	12,575	13,214	13,449	12,588
0.6569	597,053	6,562	7,073	7,437	6,576
0.6559	594,789	6,950	7,565	7,832	6,971
0.6530	592,267	8,156	8,981	9,052	8,190
0.6625	603,348	4,691	4,349	5,528	4,667

7.1.2 Linear approximation

Because utilizing the Black-Scholes formula can be intensive computationally, particularly for a large number of scenarios, it is often desirable to use a simple approximation to the formula. The simplest such approximation is to estimate changes in the option value via a linear model, which is commonly known as the "delta approximation." In this case, given an initial option value V_0 and an initial exchange rate R_0, we approximate a future option value V_1 at a future exchange rate R_1 by

[7.8] $$V_1 = V_0 + \delta (P_1 - P)$$

where

[7.9] $$\delta = \frac{\partial}{\partial R} V(P, S, \tau) \big|_{P_0}$$

is the first derivative of the option price with respect to the spot exchange rate.

For our example, V_0 is USD/DEM 0.0105 and R_0 is USD/DEM 0.65. (To compute the price of our particular option contract, we multiply V_0 by DEM 1M, the notional amount of the contract.) The value of δ for our option is –0.4919. Table 7.2 illustrates the results of the delta approximation for valuing the option's price. Note that for the delta approximation, it is still possible to utilize the standard RiskMetrics methodology without resorting to simulations.

7.1.3 Higher order approximations

It can be seen in Table 7.2 that the delta approximation is reasonably accurate for scenarios where the exchange rate does not change significantly, but less so in the more extreme cases. It is possible to improve this approximation by including the "gamma effect", which accounts for second order effects of changes in the spot rate, and the “theta effect”, which accounts for changes in the time to maturity of the option. The two formulas associated with these added effects are

[7.10] $$V_1 = V_0 + \delta(P_1 - P_0) + \frac{1}{2}\Gamma(P_1 - P_0)^2 \text{ and}$$

[7.11] $$V_1 = V_0 + \delta(P_1 - P_0) + \frac{1}{2}\Gamma(P_1 - P_0)^2 - \theta t$$

where t is the length of the forecast horizon and Γ and θ are defined by

[7.12a] $$\Gamma = \frac{\partial^2}{\partial R^2} V(P, S, \tau)\big|_{P_0} \text{ and}$$

[7.12b] $$\theta = \frac{\partial}{\partial \tau} V(P, S, \tau)\big|_{\tau_0}$$

Using the values Γ = DEM/USD 15.14 and θ = USD/DEM 0.0629 per year, we value our portfolio for each of our scenarios. The results of these approximations are displayed in Table 7.2. A plot illustrating the differences between the various methods of valuation is displayed in Chart 7.4; the delta/gamma/theta approximation is not plotted since for the values considered, it almost perfectly duplicates the full valuation case. Note that even for these higher order approximations, analytical methods exist for computing percentiles of the portfolio distribution. See, for example, the method outlined in Chapter 6.

Chart 7.4
Value of put option on USD/DEM
strike = 0.65 USD/DEM; Value in USD/DEM

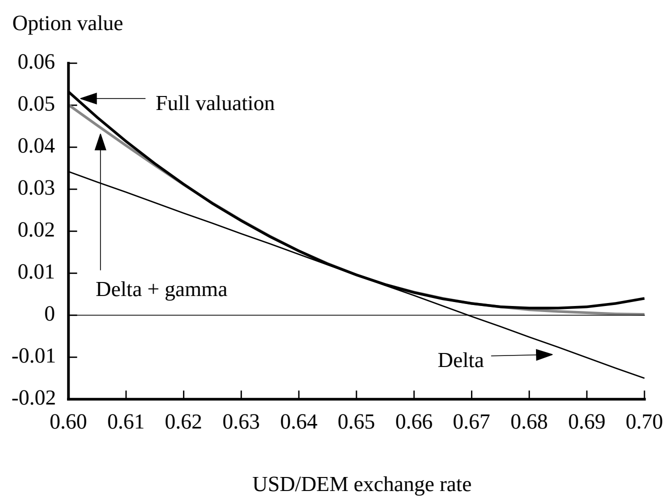

7.2 Summary

Finally, after generating a large number of scenarios and valuing the portfolio under each of them, it is necessary to make some conclusions based on the results. Clearly, one measure which we would like to report is the portfolio’s Value-at-Risk. This is done simply by ordering the portfolio return scenarios and picking out the result corresponding to the desired confidence level.

For example, to compute the 5% worst case loss using 1000 trials, we order the scenarios and choose the 50th (5% × 1000) lowest absolute return. The percentiles computed for our example under the various methods for portfolio valuation are reported in Table 7.3.

Table 7.3
Value-at-Risk for example portfolio
1000 trials

	Portfolio return (USD)			
Percentile, %	Full	Delta	Delta/Gamma	Delta/Gamma/Theta
1.0	(5,750)	(5,949)	(4,945)	(5,806)
2.5	(5,079)	(5,006)	(4,245)	(5,106)
5.0	(4,559)	(4,392)	(3,708)	(4,569)
10.0	(3,662)	(3,299)	(2,825)	(3,686)
25.0	(2,496)	(1,808)	(1,606)	(2,467)
50.0	(840)	(22)	50	(812)
75.0	915	1,689	1,813	951
90.0	2,801	3,215	3,666	2,805
95.0	4,311	4,331	5,165	4,304
97.5	5,509	5,317	6,350	5,489
99.0	6,652	6,224	7,489	6,628

Thus, at the 5% confidence level and in the full valuation case, we obtain a Value-at-Risk of USD 4,559, or about 0.75% of the current portfolio value.

A particularly nice feature of the Monte Carlo approach is that we obtain an estimate for the entire distribution of portfolio returns. This allows us to compute other risk measures if we desire, and also to examine the shape of the distribution. Chart 7.5 illustrates the return distribution for our example. Note that the distribution is significantly more skewed than the distributions for the underlying assets (see Chart 7.5), which is a result of the non-linearity of the option position.

Chart 7.5
Distribution of portfolio returns
1000 trials

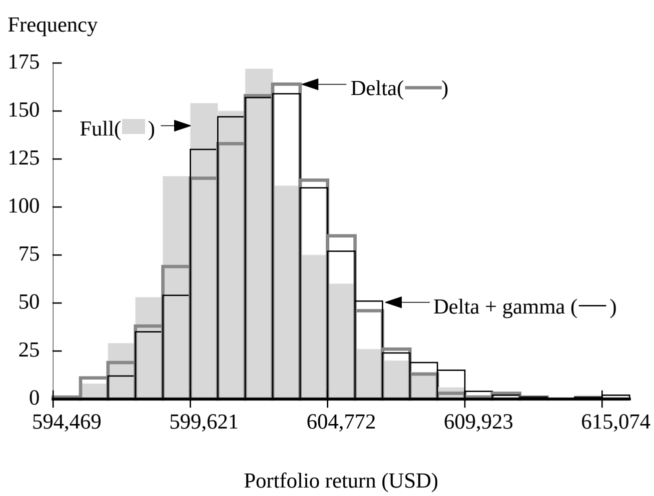

7.3 Comments

In our example, we treated the bond by assuming a lognormal process for its *price*. While this is convenient computationally, it can lead to unrealistic results since the model does not insure positive discount rates. In this case, it is possible to generate a scenario where the individual bond prices are realistic, but where the forward rate implied by the two simulated prices is negative.

We have examined a number of methods to rectify this problem, including decomposing yield curve moves into principal components. In the end, we have concluded that since regularly observed bond prices and volatilities make the problems above quite rare, and since the methods we have investigated only improve the situation slightly, it is not worth the effort to implement a more sophisticated method than what we have outlined in this chapter. We suggest a straight Monte Carlo simulation with our methodology coupled with a check for unrealistic discount or forward rates. Scenarios which yield these unrealistic rates should be rejected from consideration.