Chapter 1

Value-at-Risk

1.1. HISTORY

The term "value-at-risk" (VaR) did not enter the financial lexicon until the early 1990s, but the origins of VaR measures go further back. These can be traced to capital requirements for US securities firms of the early 20$^{\text{th}}$ century, starting with an informal capital test the New York Stock Exchange (NYSE) first applied to member firms around 1922.

Regulatory VaR Measures

The original NYSE rule[1] required firms to hold capital equal to 10% of assets comprising proprietary positions and customer receivables. By 1929, this had developed into a requirement that firms hold capital equal to:

- 5% of customer debits;
- 10% (minimum) on proprietary holdings in government of municipal bonds;
- 30% on proprietary holdings in other liquid securities; and
- 100% on proprietary holdings in all other securities.

Over time, regulators took over responsibility for setting capital requirements. In 1975, the US Securities and Exchange Commission (SEC) established a Uniform Net Capital Rule (UNCR) for US broker-dealers trading non-exempt securities. This included a system of "haircuts" that were applied to a firm's capital as a safeguard against market losses that might arise during the time it would take to

[1] See Dale (1996), pp. 60–61 and Molinari and Kibler (1983), footnote 41.

liquidate positions. Financial assets were divided into 12 categories such as government debt, corporate debt, convertible securities, and preferred stock. Some of these were further broken down into subcategories primarily according to maturity. To reflect hedging effects, long and short positions were netted within subcategories, but only limited netting was permitted across subcategories. An additional haircut was applied to any concentrated position in a single asset.

Volatility in US interest rates motivated the SEC to update these haircuts in 1980. The new haircuts were based upon a statistical analysis of historical market data. They were intended to reflect a .95-quantile of the amount of money a firm might lose over a 1-month liquidation period.[2] Although crude, the SEC's system of haircuts was a VaR measure.

Later, additional regulatory VaR measures were implemented for banks or securities firms, including:

- the UK Securities and Futures Authority 1992 "portfolio" VaR measure;
- Europe's 1993 Capital Adequacy Directive (CAD) "building-block" VaR measure; and
- the Basle Committee's[3] 1996 VaR measure based largely upon the CAD building-block measure.

In 1996, the Basle Committee approved the limited use of proprietary VaR measures for calculating bank capital requirements. In this and other ways, regulatory initiatives helped motivate the development of proprietary VaR measures.

Proprietary VaR Measures

Tracing the historical development of proprietary VaR measures is difficult because they were used by firms for internal purposes. They were not published and were rarely mentioned in the literature. One interesting document is a letter from Stephen C. Francis (1985) of Fischer, Francis, Trees & Watts to the Federal Reserve Bank of New York. He indicates that their proprietary VaR measure was similar to the SEC's UNCR but employed more asset categories, including 27 categories for cash market US Treasuries alone. He notes:

[2] See Dale (1996), p. 78.

[3] The Basle Committee on Banking Supervision is a standing committee comprising representatives from central banks and regulatory authorities. Over time, the focus of the committee has evolved, embracing initiatives designed to define roles of regulators in cross-jurisdictional situations; ensure that international banks or bank holding companies do not escape comprehensive supervision by some "home" regulatory authority; and promote uniform capital requirements so banks from different countries may compete with one another on a "level playing field." Although the Basle Committee's recommendations lack force of law, G-10 countries are implicitly bound to implement its recommendations as national laws.

> We find no difficulty utilizing on an essentially manual basis the larger number of categories, and indeed believe it necessary to capturing accurately our gross and net risk exposures.

Working at Bankers Trust, Garbade (1986) describes sophisticated VaR measures for fixed income markets that employed linear and principal component remappings to simplify computations. These may have been influenced by, but were different from, an internal VaR measure Bankers Trust implemented around 1983 for use with its risk-adjusted return on capital (RAROC) system of internal capital allocation.

Garbade today recollects[4] efforts within Bankers Trust to improve existing VaR measures following the stock market crash of 1987. During the crash, Treasury interest rates fell sharply while stock prices plummeted. Such correlated market moves are often observed during periods of market turmoil. They have motivated suggestions that correlations become more extreme during periods of elevated market volatility. Within Bankers Trust, there were several efforts to model this phenomena with mixed normal distributions. These comprised two joint-normal distributions. One was likely to be drawn from and had modest standard deviations and correlations. The other was less likely to be drawn from and had more extreme standard deviations and correlations. Using time-series methods available at the time, the researchers were unable to fit a reasonable model to available market data. They concluded that their inability to do so indicated a significant shortcoming of VaR measures then in use.

At about the same time, Chase Manhattan Bank was developing a Monte Carlo based VaR measure for use with its return on risk-adjusted capital (RORAC) internal capital allocation system. Citibank was implementing another VaR measure, also for capital allocation, which measured what the bank referred to as “potential loss amount” or PLA.[5]

A 1993 survey conducted by Price Waterhouse for the Group of 30[6] found that, at that time, among 80 responding derivatives dealers, 30% were using VaR to support market risk limits. Another 10% planned to do so.

Portfolio Theory

Directly or indirectly, regulatory and proprietary VaR measures were influenced by portfolio theory. Markowitz (1952) and Roy (1952) independently published VaR measures to support portfolio optimization. In 1952, processing power was inadequate to support practical use of such schemes, but Markowitz’s ideas

[4]Personal correspondence with the author.

[5]These VaR measures are described by Chew (1993).

[6]Founded in 1978, the Group of 30 is a nonprofit organization of senior executives, regulators, and academics. Through meetings and publications, it seeks to deepen understanding of international economic and financial issues. Results of the Price Waterhouse study are reported in Group of 30 (1994).

spawned work by more theoretically inclined researchers. Papers by Tobin (1958), Treynor (1961), Sharpe (1963, 1964), Lintner (1965), and Mossin (1966) contributed to the emerging portfolio theory. The VaR measures they employed were best suited for equity portfolios. There were few alternative asset categories, and applying VaR to these would have raised a number of modeling issues. Real estate cannot be marked to market with any frequency, making VaR inapplicable. Applying VaR to either debt instruments or futures contracts entails modeling term structures. Also, debt instruments raise issues of credit spreads. Futures that were traded at the time were primarily for agricultural products, which raise seasonality issues. Schrock (1971) and Dusak (1973) describe simple VaR measures for futures portfolios, but neither addresses term structure or seasonality issues.

Lietaer (1971) describes a practical VaR measure for foreign exchange risk. He wrote during the waning days of fixed exchange rates when risk manifested itself as currency devaluations. Since World War II, most currencies had devalued at some point; many had done so several times. Governments were secretive about planned devaluations, so corporations maintained ongoing hedges. Lietaer (1971) proposes a sophisticated procedure for optimizing such hedges. It incorporates a VaR measure with a variance of market value VaR metric. It assumes devaluations occur randomly, with the conditional magnitude of a devaluation being normally distributed. Computations are simplified using a modification of Sharpe's (1963) diagonal model. Lietaer's work may be the first instance of the Monte Carlo method being employed in a VaR measure.

Emergence of Risk Management

In 1990, risk management was novel. Many financial firms lacked an independent risk management function. This concept was practically unheard of in nonfinancial firms. The term "risk management" was not new. It had long been used to describe techniques for addressing property and casualty contingencies. Doherty (2000) traces such usage to the 1960s and 1970s when organizations were exploring alternatives to insurance, including:

- risk reduction through safety, quality control, and hazard education; and
- alternative risk financing, including self-insurance and captive insurance.

Such techniques, together with traditional insurance, were collectively referred to as risk management.

More recently, derivative dealers had been promoting "risk management" as the use of derivatives to hedge or customize market-risk exposures. For this reason, derivative instruments were sometimes called "risk management products."

The new "risk management" that evolved during the 1990s is different from either of the earlier forms. It tends to view derivatives as a problem as much as a solution. It focuses on reporting, oversight, and segregation of duties within organizations.

On January 30, 1992, Gerald Corrigan, President of the New York Federal Reserve, addressed the New York Bankers Association. His comments set a tone for the new risk management:[7]

> . . . the interest rate swap market now totals several trillion dollars. Given the sheer size of the market, I have to ask myself how it is possible that so many holders of fixed or variable rate obligations want to shift those obligations from one form to the other. Since I have a great deal of difficulty in answering that question, I then have to ask myself whether some of the specific purposes for which swaps are now being used may be quite at odds with an appropriately conservative view of the purpose of a swap, thereby introducing new elements of risk or distortion into the marketplace—including possible distortions to the balance sheets and income statements of financial and nonfinancial institutions alike.
>
> I hope this sounds like a warning, because it is. Off-balance sheet activities have a role, but they must be managed and controlled carefully, and they must be understood by top management as well as by traders and rocket scientists.

With concerns about derivatives increasing, Paul Volker, Chairman of the Group of 30, approached Dennis Weatherstone, Chairman of JP Morgan, and asked him to lead a study of derivatives industry practices. Weatherstone formed an international steering committee and a working group of senior managers from derivatives dealers; end users; and related legal, accounting, and academic disciplines. They produced a 68-page report, which the Group of 30 published in July 1993. Entitled *Derivatives: Practices and Principles*, it has come to be known as the *G-30 Report*. It describes then-current derivatives use by dealers and end-users. The heart of the study is a set of 20 recommendations to help dealers and end-users manage their derivatives activities. Topics addressed include:

- the role of boards and senior management,
- the implementation of independent risk management functions, and
- the various risks that derivatives transactions entail.

With regard to the market risk faced by derivatives dealers, the report recommends that portfolios be marked-to-market daily, and that risk be assessed with both VaR and stress testing. It recommends that end-users of derivatives adopt similar practices as appropriate for their own needs.

Although the *G-30 Report* focuses on derivatives, most of its recommendations are applicable to the risks associated with other traded instruments. For this reason, the report largely came to define the new risk management of the 1990s. The report is also interesting, as it seems to be the first published document to use the term "value-at-risk." Alternative names, such as "capital-at-risk" and "dollars-at-risk" were also used for a time and appeared earlier in the literature.[8]

[7]This incident is documented in Shirreff (1992). See Corrigan (1992) for a full text of the speech.

[8]The name "dollars-at-risk" appears as early as Mark (1991), and "capital-at-risk" as early as Wilson (1992).

Still, VaR remained a specialized tool known primarily to risk managers at financial firms. This changed in 1994 when JP Morgan introduced its free RiskMetrics service.

RiskMetrics

During the late 1980s, JP Morgan developed a firm-wide VaR system. This modeled several hundred key factors. A covariance matrix was updated quarterly from historical data. Each day, trading units would report by e-mail their positions' deltas with respect to each of the key factors. These were aggregated to express the combined portfolio's value as a linear polynomial of the risk factors. From this, the standard deviation of portfolio value was calculated. Various VaR metrics were employed. One of these was 1-day 95% USD VaR, which was calculated using an assumption that the portfolio's value was normally distributed.

With this VaR measure, JP Morgan replaced a cumbersome system of notional market risk limits with a simple system of VaR limits. Starting in 1990, VaR numbers were combined with P&L's in a report for each day's 4:15 PM Treasury meeting in New York. Those reports, with comments from the Treasury Group, were forwarded to Chairman Weatherstone.

One of the architects of the new VaR measure was Till Guldimann. His career with JP Morgan had positioned him to help develop and then promote the VaR measure within the firm. During the mid 1980s, he was responsible for the firm's asset/liability analysis. Working with other professionals, he developed concepts that would be used in the VaR measure. Later as chairman of the firm's market risk committee, he promoted the VaR measure internally. As fate would have it, Guldimann's next position placed him in a role to promote the VaR measure outside the firm.

In 1990 Guldimann took responsibility for Global Research, overseeing research activities to support marketing to institutional clients. In that capacity he managed an annual research conference for clients. In 1993, risk management was the conference theme. Guldimann gave the keynote address and arranged for a demonstration of JP Morgan's VaR system. The demonstration generated considerable interest. Clients asked if they might purchase or lease the system. Since JP Morgan was not a software vendor, they were disinclined to comply. Guldimann proposed an alternative. The firm would provide clients with the means to implement their own systems. JP Morgan would publish a methodology, distribute the necessary covariance matrix, and encourage software vendors to develop compatible software.

Guldimann formed a small team to develop something for the next year's research conference. The service they developed was called **RiskMetrics**.

It comprised a detailed technical document as well as a covariance matrix for several hundred key factors, which was updated daily. Both were distributed without charge over the Internet. The service was rolled out with considerable fanfare in October 1994. A public relations firm placed ads and articles in the financial press. Representatives of JP Morgan went on a multi-city tour to promote the service. Software vendors, who had received advance notice, started promoting compatible software.[9]

Publicized Losses

Timing for the release of RiskMetrics was excellent, as it came during a period of publicized financial losses. In February 1993, Japan's Showa Shell Sekiyu oil company reported a USD 1050MM loss from speculating on exchange rates. In December of that year, MG Refining and Marketing, a US subsidiary of Germany's Metallgesellschaft AG, reported a loss of USD 1300MM from failed hedging of long-dated oil supply commitments.

In 1994, there was a litany of losses. China's state sponsored CITIC conglomerate and Chile's state-owned Codelco copper corporation lost USD 40MM and USD 207MM trading metals on the London Metals Exchange (LME). US companies Gibson Greetings, Mead, Proctor & Gamble, and Air Products and Chemicals all reported losses from differential swaps transacted with Bankers Trust. Japan's Kashima Oil lost USD 1500MM speculating on exchange rates. California's Orange County announced losses from repos and other transactions that would total USD 1700MM. These are just a few of the losses publicized during 1994.

The litany continued into 1995. A notable example is Japan's Daiwa Bank. One of its US-based bond traders had secretly accumulated losses of USD 1100MM over a 10 year period. What grabbed the world's attention, though, was the dramatic failure of Britain's Barings PLC in February 1995. Nick Leeson, a young trader based at its Singapore office, lost USD 1400MM from unauthorized Nikkei futures and options positions. Barings had been founded in 1762. It had financed Britain's participation in the Napoleonic wars. It had financed America's Louisiana purchase and construction of the Erie Canal. Following its collapse, Barings was sold to Dutch bank ING for the price of one GBP.

By the mid-1990s, regulatory initiatives, concerns about OTC derivatives, the release of RiskMetrics, and publicized losses had created a flurry of interest in the new risk management and related techniques. Today, "value-at-risk" is not quite a household word, but it is familiar to most professionals working in wholesale financial, energy, and commodities markets.

[9] The above discussion of RiskMetrics is based upon Guldimann (2000), the author's own recollections, and private correspondence with Till Guldimann.

1.2. MEASURES

A **measure** is an operation that assigns a value to something. There are measures of length, temperature, mass, time, speed, strength, aptitude, etc. Assigned values are usually numbers, but can be elements of any ordered set. Shoe widths are sometimes assigned values from the ordered set {A, B, C, D, E}. We distinguish between:

- a **measure**, which is the operation that assigns the value, and
- a **metric**, which is an interpretation of the value.

A highway patrolman points a Doppler radar at an approaching automobile. The radar transmits microwaves, which are reflected off the auto and return to the radar. By comparing the wavelength of the transmitted microwaves to that of the reflected microwaves, the radar generates a number, which it displays. This entire process is a measure. An interpretation of that number—speed in miles/hour—is a metric.

Questionnaires are mailed to a diverse sample of 5000 households throughout the United States. They ask questions relating to:

1. business conditions in the household's area;
2. anticipated business conditions in 6 months;
3. job availability in the area;
4. anticipated job availability in 6 months; and
5. anticipated family income in 6 months.

Approximately 3500 households respond. Responses are seasonally adjusted. A statistical formula is applied to the set of responses to produce a number. This process is a measure. The Conference Board interprets the number as "consumer confidence," a unitless quantity. The interpretation is a metric.

Let's consider our first exercise. Solutions for all exercises are available online at http://www.contingencyanalysis.com.

Exercises

1.1 Describe a measure and corresponding metric that might be used in weather forecasting.

1.3. RISK MEASURES

Risk has two components:

- exposure, and
- uncertainty.

If we swim in shark-infested waters, we are *exposed* to bodily injury or death from a shark attack. We are *uncertain* because we don't know if we will be attacked. Being both exposed and uncertain, we face risk.

A **risk measure** is a measure that is applied to risks. A **risk metric** is an interpretation of such a measure. Risk metrics typically take one of three forms:

- those that quantify exposure;
- those that quantify uncertainty;
- those that quantify exposure and uncertainty in some combined manner.

Probability of rain is a risk metric that only quantifies uncertainty. It does not address our exposure to rain, which depends upon whether or not we have outdoor plans.

Credit exposure is a risk metric that only quantifies exposure. It indicates how much money we might lose if a counterparty were to default. It says nothing about our uncertainty as to whether or not the counterparty will default.

Risk metrics that quantify uncertainty—either alone or in combination with exposure—are usually probabilistic. Many summarize risk with a parameter of some probability distribution. Standard deviation of tomorrow's spot price of copper is a risk metric that quantifies uncertainty. It does so with a standard deviation. Average highway deaths per passenger-mile is a risk metric that quantifies uncertainty and exposure. We may interpret it as reflecting the mean of a probability distribution.

Exercises

1.2 Give an example of a situation that entails uncertainty but not exposure, and hence no risk.

1.3 Give an example of a situation that entails exposure but not uncertainty, and hence no risk.

1.4 In our example of the deaths per passenger-mile risk metric, for what random variable's probability distribution may we interpret it as reflecting a mean?

1.5 Give three examples of risk metrics that quantify financial risks. Choose one that quantifies exposure. Choose one that quantifies uncertainty. Choose one that quantifies uncertainty combined with exposure.

1.4. MARKET RISK

Business activities entail a variety of risks. For convenience, we distinguish between different categories of risk: market risk, credit risk, liquidity risk, etc. Although such categorization is convenient, it is only informal. Usage and definitions vary. Boundaries between categories are blurred. A loss due to widening credit spreads may reasonably be called a market loss or a credit loss, so market risk and credit risk overlap. Liquidity risk compounds other risks, such as market risk and credit risk. It cannot be divorced from the risks it compounds.

For our purposes, it is convenient to distinguish between market risk and business risk. **Market risk** is exposure to the uncertain market value of a portfolio.

A trader holds a portfolio of commodity forwards. She knows what its market value is today, but she is uncertain as to its market value a week from today. She faces market risk. **Business risk** is exposure to uncertainty in economic value that cannot be marked-to-market. The distinction between market risk and business risk parallels the distinction between market-value accounting and book-value accounting. Suppose a New England electricity wholesaler is long a forward contract for on-peak electricity delivered over the next 12 months. There is an active forward market for such electricity, so the contract can be marked to market daily. Daily profits and losses on the contract reflect market risk. Suppose the firm also owns a power plant with an expected useful life of 30 years. Power plants change hands infrequently, and electricity forward curves don't exist out to 30 years. The plant cannot be marked to market on a regular basis. In the absence of market values, market risk is not a meaningful notion. Uncertainty in the economic value of the power plant represents business risk.

Most risk metrics apply to a specific category of risks. There are market risk measures, credit risk measures, etc. Note that we do not categorize risk measures according to the specific operations those measures entail. We characterize them according to the risk metrics they are intended to support.

Gamma—as used by options traders—is a metric of market risk. There are various operations by which we might calculate gamma. We might:

- use a closed form solution related to the Black-Scholes formula;
- value the portfolio at three different underlier values and interpolate a quadratic polynomial; etc.

Each method defines a risk measure. We categorize them all as measures of gamma, not based upon the specific operations that define them, but simply because they all support gamma as a risk metric.

Exercises

1.6 Describe two different risk measures, both of which support duration as a risk metric.

1.5. VALUE-AT-RISK

Value-at-risk (VaR) is a category of market risk measures. As with any category of risk measures, we define VaR in terms of the risk metrics the measure supports.

Suppose a portfolio were to remain untraded for a certain period—say from the current time 0 to some future time 1. The portfolio's market value ${}^0p$ at the start of the period is known. Its market value ${}^1P$ at the end of the period is unknown. It is a random variable. As a random variable, we may assign it a probability distribution conditional upon information available at time 0. We might quantify the portfolio's market risk with some real-valued parameter of that conditional distribution.

Formally, a **VaR metric** is a real-valued function of:

- the distribution of ${}^1P$ conditional on information available at time 0; and
- the portfolio's current value ${}^0p$.

Standard deviation of ${}^1P$, conditional on information available at time 0, is a VaR metric:

$$ {}^0std({}^1P). \quad [1.1] $$

Conditional standard deviation of a portfolio's simple return ${}^1Z$ is a VaR metric:

$$ {}^0std({}^1Z) = {}^0std\left(\frac{{}^1P - {}^0p}{{}^0p}\right) = \frac{{}^0std({}^1P - {}^0p)}{{}^0p} = \frac{{}^0std({}^1P)}{{}^0p}. \quad [1.2] $$

If we define **portfolio loss** as

$$ {}^1L = {}^0p - {}^1P, \quad [1.3] $$

then the conditional standard deviation of ${}^1L$ is also a VaR metric:

$$ {}^0std({}^1L) = {}^0std({}^0p - {}^1P) = {}^0std({}^1P). \quad [1.4] $$

Quantiles of portfolio loss make intuitively appealing VaR metrics. If the conditional .95-quantile of ${}^1L$ is GBP 2.6MM, then such a portfolio can be expected to lose less than GBP 2.6MM on 19 days out of 20.

The functions that define VaR metrics can be fairly elaborate. An **expected tail loss** (ETL) VaR metric indicates a portfolio's expected loss conditional on that loss exceeding some specified quantile of loss.[10]

To fully specify a VaR metric, we must indicate three things:

- the period of time—1 day, 2 weeks, 1 month, etc.—between time 0 and time 1; this is the **horizon**;
- the function of ${}^0p$ and the conditional distribution of ${}^1P$;
- the currency in which ${}^0p$ and ${}^1P$ are denominated; this is the **base currency**.

Note that we always measure time in units equal to the length of the VaR horizon, so the VaR horizon always starts at time 0 and ends at time 1. We adopt a convention for naming VaR metrics:

1. The metric's name is given as the horizon, function, and currency, in that order, followed by "VaR."

[10]See Dowd (2002) for more on ETL metrics.

2. If the horizon is expressed in days without qualification, these are understood to be trading days.
3. If the function is a quantile of loss, it is indicated simply as a percentage.

For example, we may speak of a portfolio's:

- 1-day standard deviation of simple return USD VaR,
- 2-week 95% JPY VaR, or
- 1-week 90% ETL GBP VaR, etc.

Recall that risk measures are categorized according to the metrics they support. Having defined VaR metrics, we define VaR as the category of risk measures that support VaR metrics. If a risk measure supports a metric that is a VaR metric, then the measure is a **VaR measure**. If we apply a VaR measure to a portfolio, the value obtained is called a **VaR measurement** or, less precisely, the **portfolio's VaR**.

A VaR measure is just an operation—some set of computations—designed to support a VaR metric. To design a VaR measure, we generally have some financial model in mind. Models take many forms, embracing certain assumptions and drawing on fields such as portfolio theory, financial engineering, or time series analysis. Such models are the assumptions and logic that motivate a VaR measure. We call them **VaR models**.

Finally, to use a VaR measure, we must implement it. We must secure necessary inputs, code the measure as software, and install the software on computers and related hardware. The result is a **VaR implementation**.

Exercises

1.7 Which of the following represent VaR metrics:
  1. conditional variance of a portfolio's USD market value 1 week from today;
  2. conditional standard deviation of a portfolio's JPY net cash flow over the next month.
  3. beta, as defined by Sharpe's (1964) Capital Asset Pricing Model, conditional on information available at time 0.

1.8 Using the naming convention indicated in the text, name the following VaR metrics:
  1. the conditional standard deviation of a portfolio's market value, measured in AUD, 1 week from today;
  2. the conditional standard deviation of a portfolio's USD simple return over the next 3 trading days;
  3. the conditional 99% quantile of a portfolio's loss, measured in GBP, over the next day.

1.9 Consider a 1-day standard deviation of simple return JPY VaR metric. A portfolio's return is a unitless quantity; so is its conditional standard deviation of return. Must we specify a base currency (JPY) for this VaR metric? Couldn't we just call it a 1-day standard deviation of simple return VaR metric?

1.6. RISK LIMITS

VaR measures are used for a variety of tasks, including oversight, capital calculations, and portfolio optimization. The quintessential application is VaR limits.

Risk limits are a device for authorizing specific forms of risk taking. A pension fund hires an outside investment manager to invest some of its assets in intermediate corporate bonds. The fund wants the manager to take risk on its behalf, but it has a specific form of risk in mind. It doesn't want the manager investing in equities, precious metals, or cocoa futures. It communicates its intentions with contractually binding investment guidelines. These specify acceptable investments. They also specify risk limits, such as requirements that:

- the portfolio's duration always be less than 7 years;
- all bonds have a credit rating of BBB or better.

The first is an example of a market risk limit; the second of a credit risk limit.

When an organization authorizes a risk limit for risk-taking activities, it must specify three things:

1. a risk metric,
2. a risk measure that supports the metric, and
3. the limit—a value for the risk metric that is not to be exceeded.

At any point in time, a limit's **utilization** is the actual amount of risk being taken, as quantified by the risk measure. Any instance where utilization exceeds the risk limit is called a **limit violation**.

A bank's corporate lending department is authorized to lend to a specific counterparty subject to a credit exposure limit of GBP 10MM. For this purpose, the bank measures credit exposure as the sum amount of outstanding loans and loan commitments to the counterparty. The lending department lends the counterparty GBP 8MM, causing its utilization of the limit to be GBP 8MM. Since the limit is GBP 10MM, the lending department has remaining authority to lend up to GBP 2MM to the counterparty.

A metals trading firm authorizes a trader to take gold price risk subject to a 2000 troy ounce delta limit. Using a specified measure of delta, his portfolio's delta is calculated at 4:30 PM each trading day. Utilization is calculated as the absolute value of the portfolio's delta.

Market Risk Limits

For monitoring market risk, many organizations segment portfolios in some manner. They may do so by trader and trading desk. Commodities trading firms may do so by delivery point and geographic region. A hierarchy of market risk limits is typically specified to parallel such segmentation, with each segment of the portfolio having its own limits. Limits generally increase in size as you move up the hierarchy—from traders to desks to the overall portfolio; or from individual delivery points to geographic regions to the overall portfolio.

Exhibit 1.1 illustrates how a hierarchy of market risk limits might be implemented for a trading unit. A risk metric is selected, and risk limits are specified based upon this. Each limit is depicted with a cylinder. The height of the cylinder corresponds to the size of the limit. The trading unit has three trading desks, each with its own limit. There are also limits for individual traders, but only those for trading desk A are shown. The extent to which each cylinder is shaded black corresponds to the utilization of that limit. Trader A3 is utilizing almost all his limit. Trader A4 is utilizing little of hers.

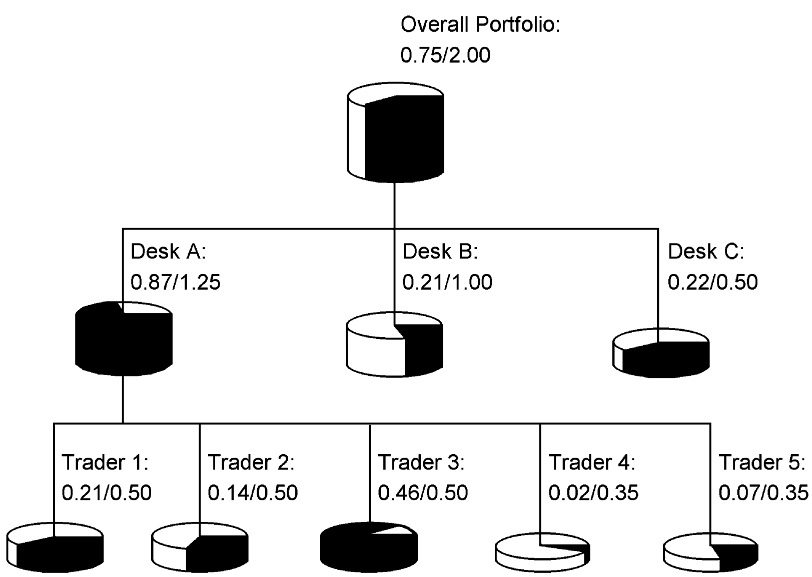

Exhibit 1.1 A hierarchy of market risk limits is illustrated for a hypothetical trading unit. A risk metric—VaR, delta, etc.—is chosen. Risk limits are specified for the portfolio and sub-portfolios based upon this. The limits are depicted with cylinders. The height of each cylinder corresponds to the size of the limit. The degree to which it is shaded black indicates current utilization of the limit. Fractions next to each cylinder indicate utilization and limit size. Units are not indicated here, as these will depend upon the particular risk metric used. Individual traders have limits, but only those for traders on desk A are indicated in the exhibit.

For such a hierarchy of risk limits to work, an organization must have a suitable risk measure to calculate utilization of each risk limit on an ongoing basis. Below, we describe three types of market risk limits, culminating with VaR limits.

Stop-Loss Limits

A stop-loss limit indicates an amount of money that a portfolio's single-period market loss should not exceed. Various periods may be used, and sometimes multiple stop-loss limits are specified for different periods. A trader might be given the following stop-loss limits:

- 1-day EUR 0.5MM,
- 1-week EUR 1.0MM,
- 1-month EUR 3.0MM.

A limit violation occurs whenever a portfolio's single-period market loss exceeds a stop-loss limit. In such an event, a trader is usually required to hedge material exposures—hence the name **stop-loss limit**.

Stop-loss limits have shortcomings. Single-period market loss is a retrospective risk metric. It only indicates risk after the financial consequences of that risk have been realized. Also, it provides an inconsistent indication of risk. If a portfolio suffers a large loss over a given period, this is a clear indication of risk. If the portfolio does not suffer a large loss, this does not indicate an absence of risk! Another problem is that traders cannot control the specific losses they incur, so it is difficult to hold traders accountable for isolated stop-loss limit violations. However, the existence of stop-loss limits does motivate traders to manage portfolios in such a manner as to avoid limit violations.

Despite their shortcomings, stop-loss limits are simple and convenient to use. Non-specialists easily understand stop-loss limits. A single risk metric can be applied consistently across an entire hierarchy of limits. Calculating utilization is as simple as marking a portfolio to market. Finally, because portfolio loss encompasses all sources of market risk, just one or a handful of limits is required for each portfolio or sub-portfolio. For these reasons, stop-loss limits are widely implemented by trading organizations.

Exposure Limits

Exposure limits are limits based upon an exposure risk metric. For limiting market risk, common metrics include: duration, convexity, delta, gamma, and vega. Crude exposure limits may also be based upon notional amounts. These are called **notional limit**s. Many exposure metrics can take on positive or negative values, so utilization may be defined as the absolute value of exposure.

Exposure limits address many of the shortcomings of stop-loss limits. They are prospective. Exposure limits indicate risk before its financial consequences are realized. Also, exposure metrics provide a reasonably consistent indication of risk. For the most part, traders can be held accountable for exposure limit violations because they largely control their portfolio's exposures.There are rare exceptions. A sudden market rise may cause a positive-gamma portfolio's delta to increase, resulting in an unintended delta limit violation.

For the most part, utilization of exposure limits is easy to calculate. There may be analytic formulas for certain exposure metrics. At worst, a portfolio must be valued under multiple market scenarios with some form of interpolation applied to assess exposure.

Exposure limits pose a number of problems. A hierarchy of exposure limits will depend upon numerous risk metrics. Not only is delta different from gamma, but crude oil delta is different from natural gas delta. Because a portfolio or subportfolio can have multiple exposures, it will require multiple exposure limits. An equity derivatives trader might have delta, gamma, and vega limits for each of 1000 equities—for a total of 3000 exposure limits.

Exposure limits are ineffective in contexts where spread trading, cross-hedging, or similar strategies minimize risk by taking offsetting positions in correlated assets. Large exposure limits are required in order to accommodate each of the offsetting positions. Because they cannot ensure reasonable hedging, the exposure limits allow for net risk far in excess of that required by the intended hedging strategy.

With the exception of notional limits, non-specialists do not easily understand exposure limits. It is difficult to know what might be a reasonable delta limit for an electricity trading desk if you don't have both:

- a technical understanding of what delta means, and
- practical familiarity with the typical size of market fluctuations in the electricity market.

This, and the sheer number of exposure limits that are often required, makes it difficult for managers to establish effective hierarchies of exposure limits.

VaR Limits

VaR limits combine many of the advantages of exposure limits and stop-loss limits. Like exposure metrics, VaR metrics are prospective. They indicate risk before its economic consequences are realized. Also like exposure metrics, VaR metrics provide a reasonably consistent indication of risk. Finally, as long as utilization is calculated for traders in a timely and ongoing manner, it is reasonable to hold them accountable for limit violations. As with exposure limits, there are rare exceptions. Consider a trader with a negative gamma position. While she is responsible for hedging the position on an ongoing basis, it is possible that a sudden move in the underlier will cause an unanticipated spike in VaR.

As with stop-loss limits, non-specialists intuitively understand VaR metrics. If a portfolio has 1-day 90% USD VaR of 7.5MM, a non-specialist understands that such a portfolio will lose less than USD 7.5MM an average of 9 days out of 10.

With VaR limits, a single metric, such as 1-day 99% USD VaR, can be applied consistently across an entire hierarchy of limits. In theory, VaR encompasses all sources of market risk. Just one limit is required for each portfolio or sub-portfolio.

VaR aggregates across assets. Depending upon the sophistication of a VaR measure, it can reflect even the most complex hedging or diversification effects. Accordingly, VaR limits are perfect for limiting risk with spread trading, cross-hedging, or similar trading strategies.

VaR limits have one significant drawback: utilization may be computationally expensive to calculate. For many portfolios, VaR is easy to calculate. It can often be done in real time on a single processor. For other portfolios, it may take minutes or hours to calculate, even with parallel processors.

Comparison

Exhibit 1.2 summarizes the strengths and weakness of stop-loss, exposure, and VaR limits. VaR limits are attractive in almost all respects. Their only significant drawback is the computational expense of calculating VaR for certain portfolios.

Characteristic	Stop-loss Limits	Exposure Limits	VaR Limits
A single metric applies across a hierarchy of limits.	✓		✓
One or few limits required per portfolio or sub-portfolio.	✓		✓
Can aggregate across exposures.	✓		✓
Easily understood by non-specialists.	✓		✓
Addresses risk prospectively.		✓	✓
Utilization provides a consistent indication of risk.		✓	✓
Traders can be held accountable for limit violations.		✓	✓
Utilization is easy to calculate.	✓	✓	

Exhibit 1.2 Characteristics of stop-loss, exposure, and VaR limits are compared. See the text for clarifications of specific issues.

1.7. EXAMPLES

Let's consider some examples of VaR measures. These will introduce basic concepts and standard notation. They will also illustrate a framework for thinking about VaR measures, which we shall formalize in Section 1.8.

The Leavens VaR Measure

Leavens (1945) published a paper describing the benefits of diversification. He accompanied his explanations with a simple example. This is the earliest known published VaR measure.

Measure time t in appropriate units. Let time $t = 0$ be the current time. Leavens considers a portfolio of 10 bonds over some horizon [0, 1]. Each bond will either mature at time 1 for USD 1000 or default and be worthless. Events of default are assumed independent. Accordingly, the portfolio's market value ${}^1P$ at time 1 is given by

$$ {}^1P = \sum_{i=1}^{10} {}^1S_i \qquad [1.5] $$

where the ${}^1S_i$ represent the individual bonds' accumulated values at time 1. Let's express this relationship in matrix notation. Let ${}^1\boldsymbol{S}$ be a random vector with components ${}^1S_i$. Let $\boldsymbol{\omega}$ be a row vector whose components are the portfolio's holdings in each bond. Since the portfolio holds one of each, $\boldsymbol{\omega}$ has a particularly simple form:

$$ \boldsymbol{\omega} = (1 \quad 1 \quad 1 \quad 1 \quad 1 \quad 1 \quad 1 \quad 1 \quad 1 \quad 1). \qquad [1.6] $$

With this matrix notation, [1.5] becomes the product:

$$ {}^1P = \boldsymbol{\omega}\,{}^1\boldsymbol{S}. \qquad [1.7] $$

Let ${}^{1|0}\phi_i$ denote the probability function, conditional on information available at time 0, of the i^{th} bond's value at time 1:

$$ {}^{1|0}\phi_i\left({}^1s_i\right) = \begin{cases} 0.9 & \text{for } {}^1s_i = 1000 \\ 0.1 & \text{for } {}^1s_i = 0 \end{cases}. \qquad [1.8] $$

Measured in USD 1000s, the portfolio's value ${}^1P$ has a binomial distribution with parameters $n = 10$ and $p = 0.9$. The probability function is graphed in Exhibit 1.3:

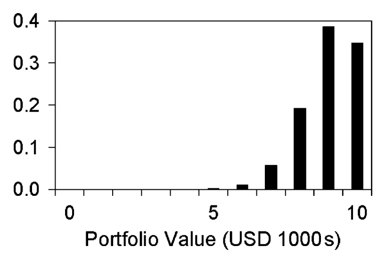

Exhibit 1.3 The market value (measured in USD 1000s) of Leavens' bond portfolio has a binomial distribution with parameters 10 and 0.9.

Writing for a non-technical audience, Leavens does not explicitly identify a VaR metric, but he speaks repeatedly of the "spread between probable losses and gains." He seems to have the standard deviation of portfolio market value in mind. Based upon this metric, his portfolio has a VaR of USD 948.69.

Some Mathematics

Our next two examples are more technical. Many readers will find them simple. Other readers—those whose mathematical background is not so strong—may find them more challenging. A note for each group:

- For the first group, the examples may tell you things you already know, but in a new way. They introduce notation and a framework for thinking about VaR that will be employed throughout the text. At points, explanations may appear more involved than the immediate problem requires. Embrace this complexity. The framework we start to develop in the examples will be invaluable in later chapters when we consider more complicated VaR measures.
- For the second group, you do not need to master the examples on a first reading. Don't think of them as a main course. They are not even an appetizer. We are taking you back into the kitchen to sample a recipe or two. Don't linger. Taste and move on. In Chapters 2 through 5, we will step back and explain the mathematics used in the examples—and used in VaR measures generally. A purpose of the examples is to provide practical motivation for those upcoming discussions.

There is a useful formula that we will use in the next two examples. We introduce it here for use in the examples, but will cover it again in more detail in Section 3.5.

Let $\boldsymbol{X}$ be a random vector with covariance matrix $\boldsymbol{\Sigma}$. Define random variable Y as a linear polynomial

$$Y = \boldsymbol{bX} + a \qquad [1.9]$$

of $\boldsymbol{X}$, where $\boldsymbol{b}$ is an n-dimensional row vector and a is a scalar. The variance of Y is given by

$$var(Y) = \boldsymbol{b\Sigma b}', \qquad [1.10]$$

where a prime $'$ indicates transposition. Formula [1.10] is a quintessential formula for describing how correlated risks combine, but there is a caveat. It only applies if Y is a linear polynomial of $\boldsymbol{X}$.

Example: Industrial Metals. Today is June 30, 2000. A US metals merchant has a portfolio of unsold physical positions in several industrial metals. We wish to calculate the portfolio's 1-week 90% USD VaR. Measure time t in weeks.

Specify the random vector

$$^{1}\boldsymbol{S} = \begin{pmatrix} ^{1}S_1 \\ ^{1}S_2 \\ ^{1}S_3 \\ ^{1}S_4 \\ ^{1}S_5 \\ ^{1}S_6 \end{pmatrix} \sim \begin{pmatrix} \text{accumulated value of a ton of aluminum} \\ \text{accumulated value of a ton of copper} \\ \text{accumulated value of a ton of lead} \\ \text{accumulated value of a ton of nickel} \\ \text{accumulated value of a ton of tin} \\ \text{accumulated value of a ton of zinc} \end{pmatrix} \qquad [1.11]$$

where accumulated values are in USD and reflect the value of a ton of metal accumulated from time 0 to time 1. Accumulated value might reflect price changes, cost of financing, warehousing, and insurance. For simplicity, we consider only price changes in this example.

Current values in USD/ton for the respective metals are

$$^{0}\boldsymbol{s} = \begin{pmatrix} ^{0}s_1 \\ ^{0}s_2 \\ ^{0}s_3 \\ ^{0}s_4 \\ ^{0}s_5 \\ ^{0}s_6 \end{pmatrix} = \begin{pmatrix} 1516.0 \\ 1719.5 \\ 476.0 \\ 7945.0 \\ 5715.0 \\ 1165.0 \end{pmatrix}. \qquad [1.12]$$

The portfolio's holdings are:

- 1000 tons of aluminum,
- 2000 tons of copper,
- 500 tons of lead,
- 250 tons of nickel,
- 1000 tons of tin, and
- 100 tons of zinc,

which we represent with a row vector:

$$\boldsymbol{\omega} = (1000 \quad 2000 \quad 500 \quad 250 \quad 1000 \quad 100). \qquad [1.13]$$

The portfolio's current value is

$$^{0}p = \boldsymbol{\omega}\,^{0}\boldsymbol{s} = 13.011\text{MM}. \qquad [1.14]$$

Its future value ^{1}P is random:

$$^{1}P = \boldsymbol{\omega}\,^{1}\boldsymbol{S}. \qquad [1.15]$$

We call this relationship a **portfolio mapping**. We represent it schematically as

$$^{1}P \xleftarrow{\boldsymbol{\omega}} {}^{1}\boldsymbol{S}. \qquad [1.16]$$

Let $^{1|0}\sigma$ and $^{1|0}\boldsymbol{\Sigma}$ be the standard deviation of ^{1}P and the covariance matrix of $^{1}\boldsymbol{S}$, both conditional on information available at time 0. Let's apply [1.10]. By

[1.15], ${}^1P$ is a linear polynomial of ${}^1\boldsymbol{S}$, so:[11]

$$^{1|0}\sigma = \sqrt{\boldsymbol{\omega}\,^{1|0}\boldsymbol{\Sigma}\boldsymbol{\omega}'}. \qquad [1.17]$$

We know $\boldsymbol{\omega}$. We need ${}^{1|0}\boldsymbol{\Sigma}$ to obtain ${}^{1|0}\sigma$. Exhibit 1.4 indicates historical metals price data.

Date	Time	Aluminum	Copper	Lead	Nickel	Tin	Zinc
	t	${}^ts_1$	${}^ts_2$	${}^ts_3$	${}^ts_4$	${}^ts_5$	${}^ts_6$
12/10/99	–29	1,516.0	1,719.5	476.0	7,945.0	5,715.0	1,165.0
12/17/99	–28	1,580.5	1,796.0	482.0	8,155.0	5,730.0	1,216.0
12/24/99	–27	1,609.0	1,834.0	474.0	8,380.0	5,700.0	1,200.0
12/30/99	–26	1,630.5	1,846.0	478.0	8,450.0	6,105.0	1,239.0
⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮
5/12/00	–7	1,455.0	1,808.5	418.0	10,040.0	5,490.0	1,170.5
5/19/00	–6	1,498.0	1,815.0	403.0	10,600.0	5,480.0	1,156.0
5/26/00	–5	1,464.0	1,793.5	432.0	10,435.0	5,405.0	1,158.0
6/2/00	–4	1,464.0	1,770.0	423.0	10,020.0	5,440.0	1,118.5
6/9/00	–3	1,456.5	1,722.5	421.0	8,480.0	5,450.0	1,099.0
6/16/00	–2	1,555.0	1,768.0	422.0	8,230.0	5,525.0	1,125.5
6/23/00	–1	1,544.5	1,767.0	416.0	7,925.0	5,515.0	1,124.0
6/30/00	0	1,564.0	1,773.5	440.5	8,245.0	5,465.0	1,148.0

Exhibit 1.4 Thirty weekly historical prices for the indicated metals. All prices are in USD per ton. Source: London Metals Exchange (LME).

Applying time-series methods described in Chapter 4, we construct

$$^{1|0}\boldsymbol{\Sigma} = \begin{pmatrix} 1{,}709 & 1{,}227 & 8 & 3{,}557 & 774 & 275 \\ 1{,}227 & 1{,}746 & 65 & 6{,}274 & 574 & 469 \\ 8 & 65 & 128 & -270 & -49 & 69 \\ 3{,}557 & 6{,}274 & -270 & 137{,}361 & -2{,}459 & 1{,}764 \\ 774 & 574 & -49 & -2{,}459 & 13{,}621 & 952 \\ 275 & 469 & 69 & 1{,}764 & 952 & 544 \end{pmatrix}. \qquad [1.18]$$

Substituting [1.13] and [1.18] into [1.17], we conclude that ${}^1P$ has conditional standard deviation ${}^{1|0}\sigma$ of 0.217MM USD.

Let ${}^{1|0}\Phi_{{}^1L}$ denote the cumulative distribution function (CDF) of portfolio loss ${}^1L$ conditional on information available at time 0. Its inverse ${}^{1|0}\Phi^{-1}_{{}^1L}$ provides quantiles of ${}^1L$. Our VaR metric is 1-week 90% USD VaR, so we seek the .90-quantile, ${}^{1|0}\Phi^{-1}_{{}^1L}(.90)$, of portfolio loss ${}^1L$.

We don't have an expression for ${}^{1|0}\Phi_{{}^1L}$. All we have is a conditional standard deviation ${}^{1|0}\sigma$ for ${}^1P$. We need additional assumptions or information. A simple solution is to assume that ${}^1P$ is conditionally normal with mean ${}^{1|0}\mu = {}^0p =$ 13.011MM. Since a normal distribution is fully specified by a mean and standard deviation, we have now specified a conditional CDF, ${}^{1|0}\Phi_{{}^1P}$, for ${}^1P$.

[11] Recall that standard deviation is the square root of variance.

The .90-quantile of portfolio loss is

$$^{1|0}\Phi^{-1}_{^1L}(.90) = {}^0p - {}^{1|0}\Phi^{-1}_{^1P}(.10). \tag{1.19}$$

A property of normal distributions is that a .10-quantile occurs 1.282 standard deviations below its mean,[12] so

$$^{1|0}\Phi^{-1}_{^1P}(.10) = {}^{1|0}\mu - 1.282\,{}^{1|0}\sigma = {}^0p - 1.282\,{}^{1|0}\sigma. \tag{1.20}$$

Substituting [1.20] into [1.19]:

$$^{1|0}\Phi^{-1}_{^1L}(.90) = 1.282\,{}^{1|0}\sigma = 0.278\text{MM}. \tag{1.21}$$

The portfolio's 1-week 90% USD VaR is USD 0.278MM. Note that ^0p dropped out of the calculations entirely, so we did not actually need to calculate its value in [1.14].

EXAMPLE: AUSTRALIAN EQUITIES. Our next example is ostensibly similar to the last. As we work through it, a number of issues will arise. These will motivate different approaches for a solution.

Today is March 9, 2000. A British trader holds a portfolio of Australian stocks. We wish to calculate the portfolio's 1-day 95% GBP VaR. The portfolio's current value ^0p is GBP 0.198MM. Let ^1P represent its value tomorrow. Define the random vector

$$^1\boldsymbol{S} = \begin{pmatrix} ^1S_1 \\ ^1S_2 \\ ^1S_3 \end{pmatrix} \sim \begin{pmatrix} \text{GBP accumulated value of a share of National Australia Bank} \\ \text{GBP accumulated value of a share of Westpac Banking Corp.} \\ \text{GBP accumulated value of a share of Goodman Fielder} \end{pmatrix}. \tag{1.22}$$

Accumulated values reflect price changes, dividends, and changes in the GBP/AUD exchange rate since time 0. The portfolio's holdings are:

- 10,000 shares of National Australia Bank,
- 30,000 shares of Westpac Banking Corp.,
- −15,000 shares of Goodman Fielder (short position),

which we represent with a row vector

$$\boldsymbol{\omega} = (10{,}000 \quad 30{,}000 \quad -15{,}000). \tag{1.23}$$

The portfolio's future value ^1P is a linear polynomial of $^1\boldsymbol{S}$:

$$^1P = \boldsymbol{\omega}\,{}^1\boldsymbol{S}. \tag{1.24}$$

We face a minor problem. In the last example, we used historical data to construct a covariance matrix for $^1\boldsymbol{S}$. In the present example, components of $^1\boldsymbol{S}$ are

[12] See Section 3.9.

denominated in GBP, but any historical data for Australian stocks will be denominated in AUD. We solve the problem with a change of variables ${}^1\boldsymbol{S} = \varphi({}^1\boldsymbol{R})$:

$$
{}^1\boldsymbol{S} = \varphi({}^1\boldsymbol{R}) = {}^1R_4 \begin{pmatrix} {}^1R_1 \\ {}^1R_2 \\ {}^1R_3 \end{pmatrix}, \tag{1.25}
$$

where

$$
{}^1\boldsymbol{R} = \begin{pmatrix} {}^1R_1 \\ {}^1R_2 \\ {}^1R_3 \\ {}^1R_4 \end{pmatrix} \sim \begin{pmatrix} \text{AUD accumulated value of a share of National Australia Bank} \\ \text{AUD accumulated value of a share of Westpac Banking Corp.} \\ \text{AUD accumulated value of a share of Goodman Fielder} \\ \text{GBP/AUD exchange rate} \end{pmatrix}. \tag{1.26}
$$

Composing $\boldsymbol{\omega}$ with φ we obtain a function $\theta = \boldsymbol{\omega} \circ \varphi$ that relates ${}^1P$ to ${}^1\boldsymbol{R}$:

$$
{}^1P = \theta({}^1\boldsymbol{R}) = {}^1R_4\left(10000\,{}^1R_1 + 30000\,{}^1R_2 - 15000\,{}^1R_3\right). \tag{1.27}
$$

This is a quadratic polynomial—the exchange rate ${}^1R_4$ combines multiplicatively with the accumulated values ${}^1R_1$, ${}^1R_2$, ${}^1R_3$. It is our portfolio mapping, and we represent it schematically as

$$
{}^1P \overbrace{\xleftarrow{\boldsymbol{\omega}} {}^1\boldsymbol{S} \xleftarrow{\varphi}}^{\theta} {}^1\boldsymbol{R}. \tag{1.28}
$$

Exhibit 1.5 provides historical data for ${}^1\boldsymbol{R}$.

Date	Time	National Australia Bank	Westpac Banking Corp.	Goodman Fielder	GBP/AUD
	t	${}^tr_1$	${}^tr_2$	${}^tr_3$	${}^tr_4$
1/10/00	−42	22.200	10.207	1.400	0.4007
1/11/00	−41	21.800	10.215	1.410	0.3990
1/12/00	−40	21.630	10.220	1.380	0.3995
1/13/00	−39	21.430	10.310	1.370	0.4057
⋮	⋮	⋮	⋮	⋮	⋮
2/29/00	−7	21.400	10.400	1.170	0.3901
3/1/00	−6	22.106	10.767	1.184	0.3828
3/2/00	−5	22.273	10.580	1.200	0.3855
3/3/00	−4	21.442	10.410	1.170	0.3847
3/6/00	−3	20.950	10.410	1.140	0.3824
3/7/00	−2	21.340	10.414	1.080	0.3826
3/8/00	−1	20.830	10.500	1.130	0.3844
3/9/00	0	20.080	10.800	1.150	0.3892

Exhibit 1.5 Two months of historical data for the GBP/AUD exchange rate and AUD prices for the indicated stocks. None of the stocks had ex-dividend dates during the period indicated. Source: Federal Reserve Bank of Chicago and Dow Jones.

Using time-series methods described in Chapter 4, we construct a conditional covariance matrix for $^1\boldsymbol{R}$:

$$^{1|0}\boldsymbol{\Sigma} = \begin{pmatrix} .156644 & .030382 & -.000135 & -.000213 \\ .030382 & .029574 & .000157 & .000053 \\ -.000135 & .000157 & .000739 & -.000010 \\ -.000213 & .000053 & -.000010 & .000015 \end{pmatrix}. \qquad [1.29]$$

Now we face another problem. We have a portfolio mapping $^1P = \theta(^1\boldsymbol{R})$ that expresses ^1P as a quadratic polynomial of $^1\boldsymbol{R}$, and we have a conditional covariance matrix $^{1|0}\boldsymbol{\Sigma}$ for $^1\boldsymbol{R}$. This is similar to the previous example where we had a portfolio mapping $^1P = \boldsymbol{\omega}\,^1\boldsymbol{S}$ that expressed ^1P as a linear polynomial of $^1\boldsymbol{S}$, and we had a covariance matrix $^{1|0}\boldsymbol{\Sigma}$ for $^1\boldsymbol{S}$. Critically, in the previous example, our portfolio mapping was linear. Now it is quadratic. In the previous example, we could apply [1.10] to obtain the conditional standard deviation of ^1P. Now we cannot.

Nonlinear portfolio mappings pose a recurring challenge for measuring VaR. There are various solutions, including:

- apply the Monte Carlo method to approximate the desired quantile;
- approximate the quadratic polynomial θ with a linear polynomial $\widetilde{\theta}$ and then apply [1.10] as before;
- assume $^1\boldsymbol{R}$ is conditionally joint-normal and apply probabilistic techniques appropriate for quadratic polynomials of joint-normal random vectors.

Each is a standard solution used frequently in VaR measures. Each has advantages and disadvantages. We will study them all in later chapters. For now, we briefly describe how each is used to calculate VaR for this Australian equities example.

EXAMPLE: AUSTRALIAN EQUITIES (MONTE CARLO TRANSFORMATION). We discuss the Monte Carlo method formally in Chapter 5. For now, an intuitive treatment will suffice. We assume $^1\boldsymbol{R}$ is joint-normal with mean vector $^{1|0}\boldsymbol{\mu} = {}^0\boldsymbol{r}$ and covariance matrix $^{1|0}\boldsymbol{\Sigma}$ given by [1.29]. Based upon these assumptions, we "randomly" generate 10,000 realizations, $^1\boldsymbol{r}^{[1]}, {}^1\boldsymbol{r}^{[2]}, \ldots, {}^1\boldsymbol{r}^{[10,000]}$, of $^1\boldsymbol{R}$. We set

$$^1p^{[k]} = \theta\left(^1\boldsymbol{r}^{[k]}\right) \qquad [1.30]$$

for each k, constructing 10,000 realizations $^1p^{[1]}, {}^1p^{[2]}, \ldots, {}^1p^{[10,000]}$ of ^1P. Results are indicated in Exhibit 1.6.

k	National Australia Bank $^1r_1^{[k]}$	Westpac Banking Corp. $^1r_2^{[k]}$	Goodman Fielder $^1r_3^{[k]}$	GBP/AUD $^1r_4^{[k]}$	Portfolio $^1p^{[k]}$
1	20.290	10.913	1.150	0.3823	196,135
2	19.392	10.333	1.153	0.3870	188,327
3	20.088	10.744	1.164	0.3917	198,117
4	20.620	11.083	1.124	0.3920	204,538
5	19.660	10.811	1.154	0.3909	196,855
6	19.973	10.806	1.162	0.3823	193,639
7	19.732	10.867	1.158	0.3902	197,437
8	19.655	10.925	1.200	0.3889	196,902
9	20.101	10.886	1.122	0.3909	199,665
10	21.136	11.064	1.129	0.3801	200,037
11	19.968	10.839	1.180	0.3881	196,804
12	20.112	10.750	1.119	0.3906	197,961
⋮	⋮	⋮	⋮	⋮	⋮
9998	20.240	10.565	1.166	0.3846	193,033
9999	19.531	10.378	1.186	0.3930	192,149
10000	20.078	11.215	1.154	0.3936	204,619

Exhibit 1.6 Results of the Monte Carlo analysis.

Realizations $^1p^{[k]}$ of ^1P are summarized with a histogram in Exhibit 1.7. We may approximate any parameter of ^1P with the corresponding sample parameter of the realizations.

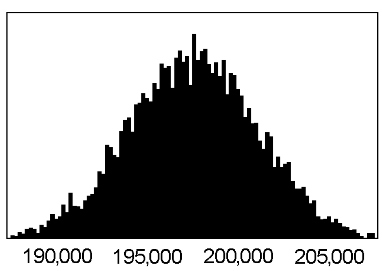

Exhibit 1.7 Histogram of realizations $^1p^{[k]}$ of the portfolio's value ^1P.

The sample .05-quantile of our realizations $^1p^{[k]}$ is USD 191,614. We use this as an approximation of the .05-quantile, $^{1|0}\Phi^{-1}_{^1P}(.05)$, of ^1P. The .95-quantile of portfolio loss is:

$$^{1|0}\Phi^{-1}_{^1L}(.95) = {}^0p - {}^{1|0}\Phi^{-1}_{^1P}(.05) \quad [1.31]$$

$$\approx 197{,}539 - 191{,}614 \quad [1.32]$$

$$= 5925. \quad [1.33]$$

The portfolio's 1-day 95% GBP VaR is approximately GBP 5925.

EXAMPLE: AUSTRALIAN EQUITIES (LINEAR REMAPPING). As an alternative solution, let's approximate θ with a linear polynomial $\widetilde{\theta}$ based upon the gradient[13] of θ. We must choose a point at which to take the gradient. A reasonable choice is ${}^0E({}^1\boldsymbol{R})$, which is the expected value of ${}^1\boldsymbol{R}$ conditional on information available at time 0. Let's assume ${}^0E({}^1\boldsymbol{R}) = {}^0\boldsymbol{r}$. We define

$$ {}^1\widetilde{P} = \widetilde{\theta}({}^1\boldsymbol{R}) = \theta({}^0\boldsymbol{r}) + \nabla\theta({}^0\boldsymbol{r})'[{}^1\boldsymbol{R} - {}^0\boldsymbol{r}]. \quad [1.34] $$

Our approximation ${}^1\widetilde{P} = \widetilde{\theta}({}^1\boldsymbol{R})$ of the portfolio mapping ${}^1P = \theta({}^1\boldsymbol{R})$ is an example of a **portfolio remapping**. We obtain ${}^0\boldsymbol{r} = (20.080, 10.800, 1.150, 0.3892)$ from Exhibit 1.5 and evaluate

$$ \theta({}^0\boldsymbol{r}) = 197{,}539 \quad [1.35] $$

and

$$ \nabla\theta({}^0\boldsymbol{r}) = \begin{pmatrix} 3{,}892 \\ 11{,}676 \\ -5{,}838 \\ 507{,}550 \end{pmatrix}. \quad [1.36] $$

Our remapping [1.34] is

$$ \begin{aligned} {}^1\widetilde{P} &= \theta({}^0\boldsymbol{r}) + \nabla\theta({}^0\boldsymbol{r})'[{}^1\boldsymbol{R} - {}^0\boldsymbol{r}] && [1.37] \\ &= \nabla\theta({}^0\boldsymbol{r})'\,{}^1\boldsymbol{R} + [\theta({}^0\boldsymbol{r}) - \nabla\theta({}^0\boldsymbol{r})'\,{}^0\boldsymbol{r}] && [1.38] \\ &= \nabla\theta({}^0\boldsymbol{r})'\,{}^1\boldsymbol{R} - 197{,}538. && [1.39] \end{aligned} $$

The portfolio remapping is represented schematically as

$$ \begin{array}{l} \overbrace{{}^1P \xleftarrow{\ \boldsymbol{\omega}\ } {}^1\boldsymbol{S} \xleftarrow{\ \varphi\ } {}^1\boldsymbol{R}}^{\theta} \\ \uparrow \\ {}^1\widetilde{P} \xleftarrow{\qquad \widetilde{\theta} \qquad} {}^1\boldsymbol{R} \end{array} \quad [1.40] $$

The upper part of the schematic is precisely schematic [1.28], indicating the original portfolio mapping ${}^1P = \theta({}^1\boldsymbol{R})$. The lower part indicates the remapping ${}^1\widetilde{P} = \widetilde{\theta}({}^1\boldsymbol{R})$. In such schematics, vertical arrows indicate approximations. ${}^1\widetilde{P}$ approximates ${}^1P$.

Because [1.39] is a linear polynomial, we can apply [1.10] to obtain the conditional standard deviation ${}^{1|0}\widetilde{\sigma}$ of ${}^1\widetilde{P}$:

$$ {}^{1|0}\widetilde{\sigma} = \sqrt{\nabla\theta({}^0\boldsymbol{r})'\,{}^{1|0}\boldsymbol{\Sigma}\,\nabla\theta({}^0\boldsymbol{r})} = 3572. \quad [1.41] $$

[13] Gradient approximations are discussed in Section 2.3.

Assume ${}^1\widetilde{P}$ is conditionally normal with this standard deviation ${}^{1|0}\widetilde{\sigma}$ and mean ${}^{1|0}\widetilde{\mu} = {}^0p$. The .05-quantile of a normal distribution occurs 1.645 standard deviations below its mean, so

$$ {}^{1|0}\Phi^{-1}_{{}^1\widetilde{P}}(.05) = {}^{1|0}\widetilde{\mu} - 1.645\,{}^{1|0}\widetilde{\sigma} = 191{,}662, \tag{1.42}$$

and the .95-quantile of portfolio loss is

$$ {}^{1|0}\Phi^{-1}_{{}^1\widetilde{L}}(.95) = {}^0p - {}^{1|0}\Phi^{-1}_{{}^1\widetilde{P}}(.05) = 5876. \tag{1.43}$$

The portfolio's 1-day 95% GBP VaR is approximately GBP 5876. This result compares favorably with our previous result of GBP 5925, which we obtained with the Monte Carlo method.

EXAMPLE: AUSTRALIAN EQUITIES (QUADRATIC TRANSFORMATION). For a third approach to calculating VaR for our Australian equities portfolio, assume that ${}^1\boldsymbol{R}$ is joint-normal with conditional mean vector ${}^0E({}^1\boldsymbol{R}) = {}^0\boldsymbol{r}$ and covariance matrix ${}^{1|0}\boldsymbol{\Sigma}$ obtained previously. Our original portfolio mapping ${}^1P = \theta({}^1\boldsymbol{R})$ defines ${}^1P$ as a quadratic polynomial of a conditionally joint-normal random vector ${}^1\boldsymbol{R}$. As we will discuss in Chapter 3, any real-valued quadratic polynomial of a joint-normal random vector can be expressed as a linear polynomial of independent normal and chi-squared random variables. In this case, the expression takes the form

$$ {}^1P = -15.02X_1 + 14.63X_2, \tag{1.44}$$

where X_1 and X_2 are independent chi-squared random variables, each with 1 degree of freedom and respective non-centrality parameters of 674.2 and 14,195.[14] This is not an approximation. The representation is exact.

There are various ways to extract a quantile of portfolio loss from a representation such as [1.44]. Two approaches that we shall discuss in Chapter 3 are:

1. approximate the desired quantile using the Cornish-Fisher (1937) expansion, and
2. invert the characteristic function of ${}^1P$ using numerical integration.

Applying the first approach to our Australian equities portfolio yields an approximate 1-day 95% GBP VaR of GBP 5854.

EXERCISES

1.10 Using a spreadsheet, extend Leavens' analysis to a bond portfolio that holds 20 bonds.
   a. Graph the resulting probability function for ${}^1P$.
   b. Based upon Leavens' "spread between probable losses and gains" VaR metric, what is the VaR of the portfolio?

[14]The chi-squared distribution is discussed in Section 3.9.

1.11 Using only the information provided in the example, which of the following VaR metrics could we evaluate for Leavens' bond portfolio:
   a. 95% quantile of loss;
   b. variance of portfolio value;
   c. standard deviation of simple return.

1.12 This exercise is based upon an equity example in Harry Markowitz's 1959 book *Portfolio Selection*. Suppose today is January 1, 1955. Measure time t in years and define:

$$ {}^{1}\boldsymbol{S} = \begin{pmatrix} {}^{1}S_1 \\ {}^{1}S_2 \\ \vdots \\ {}^{1}S_7 \end{pmatrix} \sim \begin{pmatrix} \text{accumulated value of 1 USD in AT\&T} \\ \text{accumulated value of 1 USD in American Tobacco} \\ \vdots \\ \text{accumulated value of 1 USD in Firestone} \end{pmatrix}. \qquad [1.45] $$

Each accumulated value represents the value at time 1 of an investment worth 1 USD at time 0 in the indicated stock. Accumulated values include price changes and dividends. Consider a portfolio with holdings

$$ \boldsymbol{\omega} = (10{,}000 \quad 5{,}000 \quad -1{,}000 \quad 2{,}000 \quad -5{,}000 \quad 1{,}000 \quad 6{,}000). \qquad [1.46] $$

Based upon data provided by Markowitz, we construct a conditional covariance matrix ${}^{1|0}\boldsymbol{\Sigma}$ for ${}^{1}\boldsymbol{S}$:

$$ {}^{1|0}\boldsymbol{\Sigma} = \begin{pmatrix} 0.0147 & 0.0215 & 0.0080 & 0.0145 & 0.0100 & 0.0254 & 0.0244 \\ 0.0215 & 0.0534 & 0.0162 & 0.0243 & 0.0322 & 0.0400 & 0.0490 \\ 0.0080 & 0.0162 & 0.1279 & 0.0209 & 0.0128 & 0.1015 & 0.0515 \\ 0.0145 & 0.0243 & 0.0209 & 0.0288 & 0.0113 & 0.0291 & 0.0208 \\ 0.0100 & 0.0322 & 0.0128 & 0.0113 & 0.0413 & 0.0296 & 0.0290 \\ 0.0254 & 0.0400 & 0.1015 & 0.0291 & 0.0296 & 0.1467 & 0.0900 \\ 0.0244 & 0.0490 & 0.0515 & 0.0208 & 0.0290 & 0.0900 & 0.0955 \end{pmatrix}. \qquad [1.47] $$

Calculate the portfolio's 1-year 90% USD VaR according to the following steps:
   a. Value the vector ${}^{0}\boldsymbol{s}$. (Hint: Based upon how the problem has been presented, the answer is trivial.)
   b. Using the formula ${}^{0}p = \boldsymbol{\omega}\,{}^{0}\boldsymbol{s}$, value ${}^{0}p$.
   c. Specify a portfolio mapping that defines ${}^{1}P$ as a linear polynomial of ${}^{1}\boldsymbol{S}$.
   d. Draw a schematic for your portfolio mapping.
   e. Determine the conditional standard deviation ${}^{1|0}\sigma$ of ${}^{1}P$ using [1.10].
   f. Assume ${}^{1}P$ is normally distributed with conditional mean ${}^{1|0}\mu = {}^{0}p$ and conditional standard deviation obtained in part (e). Calculate the

.10-quantile of ${}^1P$ with the formula

$$ {}^{1|0}\Phi^{-1}_{{}^1P}(.10) = {}^{1|0}\mu - 1.282\,{}^{1|0}\sigma. \qquad [1.48]$$

g. Calculate the portfolio's 1-year 90% USD VaR as

$$ {}^{1|0}\Phi^{-1}_{{}^1L}(.90) = {}^0p - {}^{1|0}\Phi^{-1}_{{}^1P}(.10). \qquad [1.49]$$

1.8. VaR MEASURES

In the previous section, we described several VaR measures. Despite a disparity in modeling techniques, our treatment was standardized. Certain concepts recurred. You are now familiar with notation such as: $\boldsymbol{\omega}$, ${}^1P$, θ, ${}^1\boldsymbol{S}$, ${}^1\boldsymbol{R}$, and ${}^{1|0}\boldsymbol{\Sigma}$.

We have many VaR measures to consider. Before long, we will stop describing entire VaR measures and start describing stand-alone components of VaR measures—much as auto enthusiasts might discuss types of brakes or fuel injectors without having a particular automobile in mind. In this sense, our discussions will have a "building block" quality. We don't want every VaR measure to be a unique monolith standing on its own. Instead, we will treat them as modular. Avoiding the top-down approach of discussing Toyotas, Fords, and Mercedes, we will take a bottom-up approach, discussing fuel injectors, suspension systems, and brakes. To this end, we must identify the essential components that make up any VaR measure. In doing so, we will lay out a framework for much of this book.

RISK FACTORS

A **risk factor** is any random variable ${}^1Q_i$ whose value will be realized during the interval (0,1] and will affect the market value of a portfolio at time 1. A **risk vector** ${}^1\boldsymbol{Q}$ is a random vector of risk factors. If a risk vector reflects a future value of some time series, we may speak of its **current value** ${}^0\boldsymbol{q}$ or **historical values** ${}^0\boldsymbol{q}$, ${}^{-1}\boldsymbol{q}$, ${}^{-2}\boldsymbol{q}$, ${}^{-3}\boldsymbol{q}$, . . .

One particular risk factor and two risk vectors play important roles in VaR measures. We give them special names and notation. These are:

- the portfolio's future value ${}^1P$;
- the asset vector ${}^1\boldsymbol{S}$; and
- the key vector ${}^1\boldsymbol{R}$.

The **portfolio's future value** ${}^1P$ represents the market value at time 1 of the portfolio for which VaR is to be measured. The portfolio is assumed fixed in the sense that it will not be traded during the period [0, 1], and no assets will be added or withdrawn. This does not preclude traders or portfolio managers from trading!

It simply means that a VaR measure quantifies the market risk of a portfolio based upon its composition at time 0. The VaR measure can recognize changes in the portfolio's composition during the period [0, 1] due to planned events such as options expiring, dividends being paid, or scheduled payments being made on a swap.

We are interested in the portfolio's current value ${}^0p$ only if a VaR metric depends upon it. We generally do not consider or attempt to define prior historical portfolio values. Asset vector ${}^1\boldsymbol{S}$ has **asset values** ${}^1S_i$ as components. These represent accumulated values of specific assets that may make up a portfolio. Realizations ${}^1s_i$ may be negative, so our definition recognizes no accounting distinction between assets and liabilities. Accumulated value is denominated in the base currency employed by the VaR metric. It may reflect such variables as capital gains, dividends, coupons, margin payments, reinvestment income, storage costs, insurance, financing, changes in exchange rates, leasing income, etc.

Mathematically, we define a **portfolio** as a pair $({}^0p, {}^1P)$, where the constant ${}^0p$ is the portfolio's current value, and the random variable ${}^1P$ is the portfolio's future value. Similarly, we mathematically define an **asset** as a pair $({}^0s_i, {}^1S_i)$, where ${}^0s_i$ is the asset's current value, and ${}^1S_i$ is the asset's future value.

We have considerable leeway in how we select what financial instruments to represent with assets. This may affect VaR results. Consider an investor who borrows EUR 100,000 and invests it in Hoechst stock. We might model the portfolio three different ways:

1. as comprising holdings in two assets whose values ${}^1S_1$ and ${}^1S_2$ represent the accumulated values of the stock and the financing, respectively;
2. as comprising a single asset whose value ${}^1S_1$ represents the accumulated value of the stock less the accumulated value of its financing;
3. as comprising a single asset whose value ${}^1S_1$ represents the accumulated value of the stock.

The first two representations are financially equivalent. One approach (probably the first) will be computationally easier to work with, but both will result in the same VaR. The third representation is different. It excludes financing from the portfolio. With it, the random variable ${}^1P$ represents something different than it does with the first two approaches.

As we shall see, every VaR measure must directly characterize a conditional probability distribution for some vector of risk factors, such as prices, interest rates, spreads, or implied volatilities. Those risk factors ${}^1R_i$ are called **key factors**. They are the components of the **key vector** ${}^1\boldsymbol{R}$. Occasionally, we use asset values ${}^1S_i$ as key factors. This was the case in our examples of Leavens' VaR measure and the VaR measure for industrial metals. We explore the role of key factors in more detail shortly.

Holdings

When we design a VaR measure, we must decide what financial assets to represent with mathematical assets (${}^0s_i$, ${}^1S_i$). We might measure equity positions in shares or round lots. In Exercise 1.12, we measured them as the number of USD held in a given stock at time 0. Positions in cocoa might be measured in pounds, bags, or tons. The choice of units is largely arbitrary, but it must be explicit if we are to define portfolio holdings.

A portfolio's **holdings** is a row vector $\boldsymbol{\omega}$ indicating the number ω_i of units held by the portfolio in each asset.

Mappings

In mathematics, a mapping is a function. The words are synonyms. In the context of VaR, we reserve the word "mapping" for functions relating specific risk vectors to one another. If ${}^1\boldsymbol{Q}$ and ${}^1\dot{\boldsymbol{Q}}$ are risk vectors, a **mapping** is a functional relationship:

$$ {}^1\boldsymbol{Q} = \varphi({}^1\dot{\boldsymbol{Q}}). \qquad [1.50] $$

We call φ the **mapping function**.

A **portfolio mapping** is a mapping that defines a portfolio's value ${}^1P$ as a function of some risk vector ${}^1\boldsymbol{Q}$:

$$ {}^1P = \varphi({}^1\boldsymbol{Q}). \qquad [1.51] $$

Portfolio mappings play a simple but inevitable role in VaR measures. Let's focus on two of our earlier examples: Leavens' VaR measure and our Australian equities VaR measures.

To calculate a portfolio's VaR, we must calculate the value of some function—VaR metric—of ${}^0p$ and the conditional distribution of ${}^1P$. We interpret ${}^1P$ as the portfolio's market value at time 1, but this is not a definition. Mathematically, there are two ways we may define the random variable ${}^1P$:

1. we can directly specify a conditional distribution for ${}^1P$;
2. we can define ${}^1P$ as a function of some random vector.

The first approach is hardly feasible. Portfolios and financial markets tend to be complicated, so it is difficult to directly specify a conditional distribution for ${}^1P$. Inevitably, we define ${}^1P$ using the second approach—which leads to portfolio mappings. Both the Leavens and Australian equities VaR measures define ${}^1P$ as a function of some asset vector ${}^1\boldsymbol{S}$:

$$ {}^1P = \boldsymbol{\omega}\,{}^1\boldsymbol{S}. \qquad [1.52] $$

We interpret ${}^1\boldsymbol{S}$ as a vector of accumulated values, but this is not a definition. To complete our definition of ${}^1P$, we must mathematically define ${}^1\boldsymbol{S}$. As with ${}^1P$, there are two ways to define ${}^1\boldsymbol{S}$:

1. we can directly specify a conditional distribution for ${}^1\boldsymbol{S}$;
2. we can define ${}^1\boldsymbol{S}$ as a function of some other random vector.

At this point, Leavens uses the first approach. He specifies a conditional distribution for ${}^1\boldsymbol{S}$ and uses this to infer a binomial distribution for ${}^1P$. We schematically represent Leavens' portfolio mapping as

$$ {}^1P \xleftarrow{\boldsymbol{\omega}} {}^1\boldsymbol{S}. \qquad [1.53]$$

The Australian equities VaR measures don't stop there. Rather than directly specify a joint distribution for ${}^1\boldsymbol{S}$, they define ${}^1\boldsymbol{S}$ as a mapping of another random vector ${}^1\boldsymbol{R}$. We schematically represent the resulting portfolio mapping as

$$ {}^1P \overbrace{\xleftarrow{\boldsymbol{\omega}} {}^1\boldsymbol{S} \xleftarrow{\varphi}}^{\theta} {}^1\boldsymbol{R}. \qquad [1.54]$$

No matter how many mappings are composed, ultimately ${}^1P$ must be defined as a function of some random vector for which we directly characterize a joint distribution. That random vector is the key vector ${}^1\boldsymbol{R}$. We denote the mapping function that relates ${}^1P$ to its key vector ${}^1\boldsymbol{R}$ with θ. Accordingly, the notation

$$ {}^1P = \theta({}^1\boldsymbol{R}) \qquad [1.55]$$

recurs frequently in this text. An exception is if asset values are used as key factors. In this case, the relationship is

$$ {}^1P = \boldsymbol{\omega}\,{}^1\boldsymbol{S}, \qquad [1.56]$$

and ${}^1\boldsymbol{S}$ plays the dual role of asset vector and key vector.

Here, we have described not only portfolio mappings, but also a general procedure for constructing them. Portfolio mappings constructed in this manner—starting with asset vector ${}^1\boldsymbol{S}$ and holdings $\boldsymbol{\omega}$, and perhaps mapping ${}^1\boldsymbol{S}$ to some key vector ${}^1\boldsymbol{R}$—are called **primary mappings**. The name distinguishes them from portfolio mappings constructed as remappings. All portfolio mappings stem from primary mappings. They either are left in that form, or are approximated using one or more remappings. We discuss primary mappings in Chapter 8.

INFERENCE

In order to characterize a distribution for ${}^1P$ conditional on information available at time 0, we must characterize a conditional distribution for ${}^1\boldsymbol{R}$. We do so with an **inference procedure**. It is not always necessary to fully specify a distribution. We require only information sufficient to value our chosen VaR metric. Some inference procedures characterize the conditional distribution of $^1\boldsymbol{R}$ with just a covariance matrix. We say an inference procedure is **complete** if it fully specifies a conditional distribution for $^1\boldsymbol{R}$. Otherwise it is **incomplete**.

Inference procedures take various forms. Leavens (1945) simply makes up a distribution suitable for his example. In practice, techniques of time series analysis are employed—in conjunction with financial theory—to obtain a reasonable characterization. We discuss inference procedures in Chapter 7.

Transformations

A **transformation procedure**—or **transformation**—characterizes a conditional distribution for ^1P and uses this characterization to value a desired VaR metric. Recall that risk comprises two components:

- exposure, and
- uncertainty.

A portfolio mapping $^1P = \theta(^1\boldsymbol{R})$ incorporates both. The characterization of a conditional distribution of $^1\boldsymbol{R}$ reflects our uncertainty. The mapping function θ reflects our exposure. The challenge for a transformation procedure is to combine both components to characterize a conditional distribution for ^1P. To intuitively understand what this entails, consider some simple examples.

A portfolio's value depends upon a single normally distributed key factor $^1R_1$. The mapping function θ is a linear polynomial. The situation is depicted in Exhibit 1.8:

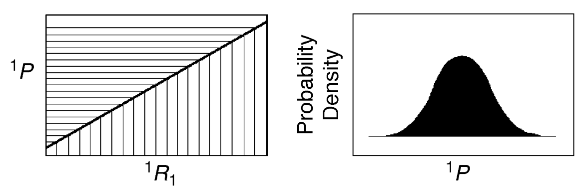

Exhibit 1.8 A linear mapping function θ is applied to a key factor $^1R_1$. This is illustrated intuitively by mapping evenly spaced realizations for $^1R_1$ through the mapping function. The output values for ^1P are also evenly spaced, indicating that the mapping function causes no distortion. Since $^1R_1$ is conditionally normal, so is ^1P.

The graph on the left depicts the mapping function θ. Evenly spaced realizations for $^1R_1$ have been mapped into corresponding realizations for ^1P. The resulting realizations of ^1P are also evenly spaced, indicating that θ imparts no distortions. Since $^1R_1$ is conditionally normal, ^1P will also be conditionally normal, as illustrated in the graph on the right.

For our second example, consider a portfolio comprising a single call option with a conditionally normal key factor ${}^1R_1$ as its underlier. To avoid a need for additional key factors, treat applicable interest rates and implied volatilities as constant.

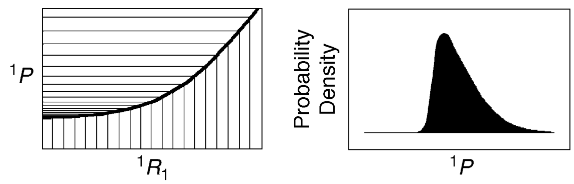

Exhibit 1.9 A nonlinear mapping function θ is applied to a conditionally normal key factor ${}^1R_1$. The result is a conditionally non-normal portfolio value ${}^1P$. This is illustrated intuitively by mapping evenly spaced realizations for ${}^1R_1$ through the mapping function. The corresponding values for ${}^1P$ are not evenly spaced, reflecting how the mapping function distorts the distribution of ${}^1P$.

In Exhibit 1.9, the left graph depicts the familiar "hockey stick" mapping function of a call option. Evenly spaced realizations for ${}^1R_1$ do not map into evenly spaced realizations for ${}^1P$, so the mapping function causes distortions. Since ${}^1R_1$ is conditionally normal, the resulting distribution of ${}^1P$ is conditionally non-normal, as illustrated on the right.

Our third example considers a long-short options position applied to a short position in the underlier. The mapping function θ, which is illustrated in the left graph of Exhibit 1.10, causes realizations of ${}^1P$ to cluster in two regions. If the underlier ${}^1R_1$ is conditionally normal, ${}^1P$ will have the dramatically non-normal conditional distribution shown on the right.

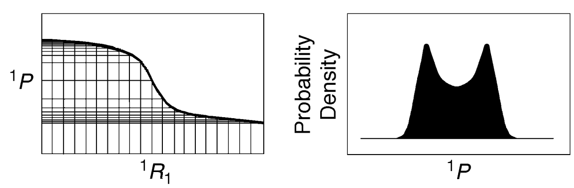

Exhibit 1.10 A long-short options position can result in a bimodal distribution for ${}^1P$.

These are simple examples, especially since each entails a single key factor. Practical VaR measures often entail 100 or more key factors. If a portfolio holds complex instruments such as exotic derivatives or mortgage-backed securities, a mapping function can be extremely complex. Such issues can make it difficult to design a practical transformation procedure.

We say a transformation procedure is **complete** if its characterization of the conditional distribution for ${}^1P$ is sufficiently general to support any practical VaR metric. Otherwise, the transformation procedure is **incomplete**. For example, if a transformation characterizes the conditional distribution of ${}^1P$ with a mean and a standard deviation, it is incomplete. If it characterizes it as conditionally normal with a specified mean and standard deviation, it is complete. We call a VaR measure **complete** or **incomplete** according to whether its transformation procedure is complete.

In our examples of Section 1.7, we illustrated three types of transformations:

1. linear transformations,
2. quadratic transformations, and
3. Monte Carlo transformations.

The first applies if a portfolio mapping function θ is a linear polynomial. The second applies if θ is a quadratic polynomial and ${}^1\boldsymbol{R}$ is joint-normal. The third applies quite generally and is one example of a category of transformations called numerical transformations. We discuss transformation procedures in Chapter 10.

Remappings

All the VaR measures we have considered so far entail modest calculations. They apply to small portfolios that are easy to value. When we develop VaR measures for real portfolios, this will change.

Every VaR measure employs—explicitly or implicitly—a primary mapping ${}^1P = \theta({}^1\boldsymbol{R})$. Primary mappings can be complicated. This occurs for two reasons:

1. The mapping function θ may be complicated—Mapping functions are formulas for marking-to-market a portfolio as of time 1. They are constructed using techniques of financial engineering. All the computational challenges that arise with financial engineering can arise with θ.

2. The key vector ${}^1\boldsymbol{R}$ may be complicated—VaR measures are sometimes implemented with 1000 or more key factors ${}^1R_i$. Also, the joint distribution of ${}^1\boldsymbol{R}$ may be difficult to work with.

Such complexity can make it difficult to directly apply a transformation procedure. This is especially true if both a primary mapping and a transformation procedure employ the Monte Carlo method—resulting in nested Monte Carlo analyses.

Consider a portfolio holding 300 exotic derivatives, each of which can only be valued using the Monte Carlo method. The primary mapping has the form

$$\overbrace{{}^1P \xleftarrow{\boldsymbol{\omega}} {}^1\boldsymbol{S} \xleftarrow{\varphi} {}^1\boldsymbol{R}}^{\theta}, \qquad [1.57]$$

where key factors ${}^1R_i$ represent values for underliers, implied volatilities, and discount factors. Valuing the mapping ${}^1\boldsymbol{S} = \varphi({}^1\boldsymbol{R})$ for a specific realization ${}^1\boldsymbol{r}^{[k]}$ requires 300 Monte Carlo analyses, one for each derivative's value: ${}^1s_i^{[k]} = \varphi_i({}^1\boldsymbol{r}^{[k]})$. Valuing a realization ${}^1p^{[k]}$ of ${}^1P$ based upon one realization ${}^1\boldsymbol{r}^{[k]}$ entails performing all 300 of these Monte Carlo analyses.

Suppose we employ a Monte Carlo transformation procedure to calculate the portfolio's VaR. This will nest the 300 valuation Monte Carlo analyses within the Monte Carlo transformation. The Monte Carlo transformation might calculate 10,000 realizations ${}^1p^{[k]}$. Since each entails 300 valuation Monte Carlo analyses, the entire analysis will entail 10,000(300) = 3,000,000 Monte Carlo analyses. This is a staggering computational load.

To make a transformation less computationally expensive, we might replace a primary mapping ${}^1P = \theta({}^1\boldsymbol{R})$ with an approximation ${}^1\widetilde{P} = \widetilde{\theta}({}^1\widetilde{\boldsymbol{R}})$, which we call a remapping. In our (second) Australian equities example, we considered a simple remapping. The above example of nested Monte Carlo analyses illustrates an extreme case. Here, a remapping would be crucial.

Formally, a remapping is an approximation of a risk vector ${}^1\boldsymbol{Q}$ with some other risk vector ${}^1\widetilde{\boldsymbol{Q}}$. We describe remappings more generally in Chapter 9. For now, we are interested in remappings ${}^1\widetilde{P}$ of ${}^1P$. If we have a portfolio mapping ${}^1P = \theta({}^1\boldsymbol{R})$, such remappings may take three forms:

1. A **function remapping** approximates ${}^1P = \theta({}^1\boldsymbol{R})$ by replacing θ with an approximate mapping function $\widetilde{\theta}$, so ${}^1\widetilde{P} = \widetilde{\theta}({}^1\boldsymbol{R})$.
2. A **variables remapping** approximates ${}^1P = \theta({}^1\boldsymbol{R})$ by replacing ${}^1\boldsymbol{R}$ with alternative key vector ${}^1\widetilde{\boldsymbol{R}}$, so ${}^1\widetilde{P} = \theta({}^1\widetilde{\boldsymbol{R}})$.
3. A **dual remapping** approximates ${}^1P = \theta({}^1\boldsymbol{R})$ by replacing both θ and ${}^1\boldsymbol{R}$, so ${}^1\widetilde{P} = \widetilde{\theta}({}^1\widetilde{\boldsymbol{R}})$.

The first and third forms are most common. Many function remappings approximate a portfolio mapping function θ with a linear or quadratic polynomial $\widetilde{\theta}$ to facilitate use of a linear or quadratic transformation. Many dual remappings replace a high-dimensional ${}^1\boldsymbol{R}$ with a lower dimensional ${}^1\widetilde{\boldsymbol{R}}$. Principal component analysis, which we discuss in Chapter 3, can be useful for this purpose. Remappings may be applied to primary mappings or to other remappings—approximating approximations.

Function and dual remappings entail a change of key factors. This raises an important issue. Key factors are specific to a portfolio. Portfolio $({}^0p, {}^1P)$ has key vector ${}^1\boldsymbol{R}$. A remapped portfolio $({}^0\widetilde{p}, {}^1\widetilde{P})$ may have the same key vector ${}^1\boldsymbol{R}$ (as is the case with a function remapping) or it may have a different key vector ${}^1\widetilde{\boldsymbol{R}}$. We discuss remappings in Chapter 9.

Summary

Recall our definition of measure from Section 1.2:

> A measure is an operation that assigns a value to something.

A VaR measure is an operation that assigns a value to a portfolio. That operation comprises various procedures, which we have defined above. Exhibit 1.11 relates these to one another in a general schematic.

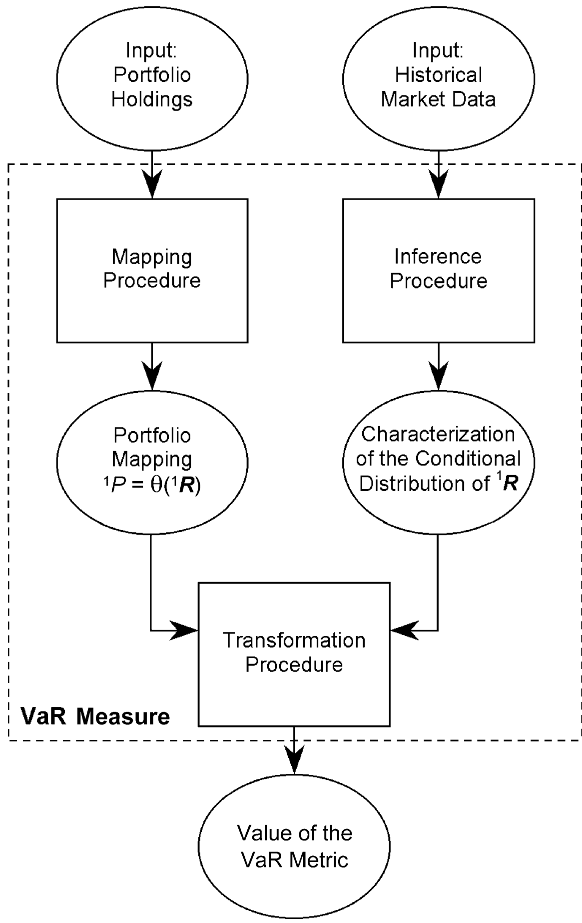

Exhibit 1.11 VaR measures follow a common general scheme, which is depicted here. Procedures are indicated with rectangles; inputs and outputs for procedures are indicated with ovals.

Specific VaR measures vary in certain respects, but all conform generally to the scheme of Exhibit 1.11. They accept both a portfolio's holdings and historical market data as inputs. A **mapping procedure** specifies a portfolio mapping function θ, which may reflect a primary portfolio mapping or a portfolio remapping. An **inference procedure** characterizes a conditional distribution for $^1\boldsymbol{R}$. It generally employs techniques of time-series analysis.

The outputs of the mapping and inference procedures reflect the two components of risk. The mapping function θ reflects exposure. The characterization of the conditional distribution of $^1\boldsymbol{R}$ reflects uncertainty. A transformation procedure combines these two components to somehow characterize the conditional distribution of ${}^1P$. It then uses that characterization to determine a value for the desired VaR metric, which is the output for the VaR measure.

To characterize the conditional distribution of ${}^1P$, the transformation procedure may employ results from probability theory as well as methods of numerical integration, such as the Monte Carlo method. The characterization may take many forms—a probability density function (PDF), a characteristic function, certain parameters of the distribution of ${}^1P$, a realization of a sample[15] from the distribution of ${}^1P$, etc. If the characterization is sufficiently general to calculate any practical VaR metric, we say the transformation is **complete**. Otherwise, it is **incomplete**. We call a VaR measure **complete** or **incomplete** according to whether or not its transformation procedure is complete.

Exercises

1.13 Below are informal descriptions of three portfolio mappings and three schematics of portfolio mappings. Match each description with the corresponding schematic.

a. Portfolio value depends upon key factors ${}^1R_i$ representing exchange rates, implied volatilities, and interest rates in various currencies.
b. A stock portfolio is modeled as a function of individual stocks' single-period returns. For simplicity, all return pairs are assumed to have the same correlation.
c. A portfolio holds options and futures on gold. Its market value is approximated as a quadratic polynomial of applicable risk factors.

Schematic 1:

$$\begin{array}{l} \overbrace{{}^1P \xleftarrow{\boldsymbol{\omega}} {}^1\boldsymbol{S} \xleftarrow{\varphi} {}^1\boldsymbol{R}}^{\theta} \\ \uparrow \\ {}^1\tilde{P} \xleftarrow{\qquad\tilde{\theta}\qquad} {}^1\boldsymbol{R} \end{array} \qquad [1.58]$$

Schematic 2:

$$\overbrace{{}^1P \xleftarrow{\boldsymbol{\omega}} {}^1\boldsymbol{S} \xleftarrow{\varphi} {}^1\boldsymbol{R}}^{\theta} \qquad [1.59]$$

[15] As obtained with a Monte Carlo transformation.

Schematic 3:

$$\begin{array}{c}
\overbrace{{}^1P \xleftarrow{\boldsymbol{\omega}} {}^1\boldsymbol{S} \xleftarrow{\varphi} {}^1\boldsymbol{R}}^{\theta} \\
\uparrow \qquad\quad \uparrow \qquad\quad \uparrow \\
{}^1\widetilde{P} \xleftarrow{\boldsymbol{\omega}} {}^1\widetilde{\boldsymbol{S}} \xleftarrow{\varphi} {}^1\widetilde{\boldsymbol{R}}
\end{array} \qquad [1.60]$$

1.14 Exhibit 1.12 illustrates three portfolio mapping functions θ for portfolios whose values ${}^1P$ depend upon a single key factor ${}^1R_1$. As we did in Exhibits 1.8, 1.9, and 1.10, sketch what each conditional PDF of ${}^1P$ might look like assuming ${}^1R_1$ is conditionally normal with its mean at the mid-point of each graph.

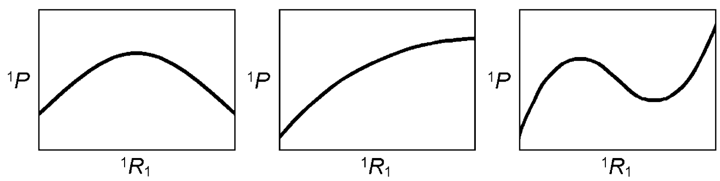

Exhibit 1.12 Portfolio mapping functions θ for Exercise 1.14.

1.15 Describe portfolios whose mapping functions might appear like those of the previous exercise.

1.9. FURTHER READING

For more information on the origins of regulatory VaR measures, see Dale (1996) and Molinari and Kibler (1983). Bernstein (1992) and Markowitz (1999) describe the origins of VaR measures in the context of portfolio theory. See Guldimann (2000) for the history of RiskMetrics. Our discussion of measures is largely operational; see Lad (1996). The classic text—and still one of the best—on financial risk management is the Group of 30 (1993) report. See also Culp (2001). For market risk management specifically, see Goldman Sachs and SBC Warburg Dillon Read (1998). For an alternative treatment of VaR measures, see Morgan Guaranty (1996), Dowd (2002), or Marrison (2002).