COLUMBIA
BUSINESS
SCHOOL

Value at Risk (A)

The collapse of Barings Bank, the widely publicized derivatives losses of Orange County and Metallgesellschaft Refining and Manufacturing, the near-demise of Long Term Capital Management, and numerous other related incidents have focused the attention of regulators and financial institutions on improved methods for measuring and managing financial risks. The increasing complexity and use of derivative securities to repackage and redistribute risks have rendered inadequate traditional accounting-based measures to determine the capital required to protect against trading losses. In just the past few years, *value at risk* (VaR) has become the most important benchmark for measuring risk in portfolios of diverse and often complex instruments.

The following excerpt from the 1998 Chase annual report is typical of the way financial institutions use and measure VaR:

> *Chase's two principal risk measurement tools are VaR and stress testing. VaR measures risk in an everyday environment, while stress testing measures market risk in an abnormal market environment. The VaR, a dollar amount, is a forward looking estimate of the potential for loss. The VaR looks forward one trading day, and is calculated as the loss level expected to be exceeded with a 1 in 100 chance.*

A 1996 Goldman Sachs research report includes the following description:

> *Value at risk is a measure of a point in the distribution of possible outcomes. It has two parameters: a horizon and a probability. For example, a common regulatory definition of VaR is the amount of capital that you should expect to lose no more than once in a hundred two-week intervals, given your current*

positions. At Goldman Sachs, we commonly focus on an amount of capital that we should expect to lose no more than once per year in a given day. We think of this not as a "worst case," but rather as a regularly occurring event with which we should be comfortable.[2]

In a joint report[3] The Department of the Treasury, the Federal Reserve, and the Federal Deposit Insurance Corporation offer the following description:

> *The VaR measure represents an estimate of the amount by which an institution's position in a risk category could decline due to general market movements during a given holding period.*

This US report follows an international accord adopted by the Group of Ten countries[4] through the Basle Committee on Banking Supervision. A 1995 report of the Basle Committee was instrumental in focusing attention on VaR as a measure of risk.

Among the more precise definitions of VaR is the following one from J.P. Morgan's 1996 *RiskMetrics Technical Document*:

> *Value at risk is a measure of the maximum potential change in value of a portfolio of financial instruments over a pre-set horizon. VaR answers the question: how much can I lose with x% probability over a given time horizon.*

To put it even more plainly, VaR is a percentile of the profit and loss distribution of a portfolio over a specified horizon. A 95% VaR is the size of the loss that will be exceeded with only 5% probability; a 99% VaR is a loss that will be exceeded with only 1% probability. To complete the specification, we need to indicate a time horizon — one day and ten days are commonly used. If we say that a portfolio has a 95% one-day VaR of \$100 million, we mean that there is only a 5% chance that the portfolio will lose more than \$100 million over the next day. The quotes above indicate that both Chase and Goldman Sachs look at one-day VaR; Chase uses a loss probability of 1% whereas Goldman Sachs uses $1/250 = 0.4\%$, assuming 250 business days in a year.

Figure 1 shows a hypothetical profit and loss (P&L) distribution for a portfolio over some time horizon (e.g., one day). The horizontal axis

[2] R. Litterman, "Hot Spots and Best Hedges," Risk Management Series, Goldman Sachs, New York (1996)

[3] *Federal Register*, vol. 61, no. 174, September 6, 1996.

[4] The G-10 countries are Belgium, Canada, France, Germany, Italy, Japan, the Netherlands, Sweden, Switzerland, the United Kingdom, and the United States.

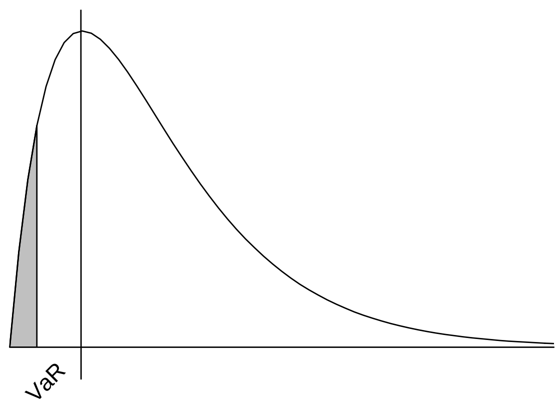

Figure 1: The value at risk for a hypothetical P&L distribution. The VaR is a lower percentile of the distribution. The shaded area corresponds to 5% of the area under the distribution, so the indicated point gives a 95% VaR.

corresponds to different levels of profit or loss; the vertical axis gives the relative probabilities of the various outcomes. The shaded area corresponds to 5% of the area under the curve. Thus, there is a 5% chance that the loss over the time horizon will exceed the VaR.

Capital Adequacy Requirements

Financial institutions have many reasons for measuring portfolio risk; regulators, on the other hand, are primarily concerned with the solvency of the financial system. To preserve solvency, regulators require that financial firms hold adequate capital to sustain trading losses. To the firms, holding capital represents a cost, so they would often prefer to be able to hold less capital.

Historically, capital requirements have been based on a rather crude view of risk: a fixed percentage reserve requirement was assigned to each type of asset. For example, certain types of loans might require reserves in the amount of 2% of the principal and others might entail an 8% capital charge. Summing the charges over all assets yields the total capital required.

This “building-block” approach ignores the possibility of risk reduction through diversification. It also ignores the possibility of changing market conditions. It may be well-suited to a bank dealing only in standard loans, but it is entirely inadequate for institutions with large positions in swaps, options, and other derivative securities.

The impetus for rethinking capital requirements has come in part from the financial institutions themselves. With the increasing globalization of financial services, banks based in countries with strict regulations sometimes find themselves at a competitive disadvantage compared with banks based in countries with lax requirements. Setting international standards through bodies like the Basle Committee is in part an attempt to level the playing field across borders.

By basing capital requirements on VaR, regulatory agencies address some of the shortcomings of the building-block approach.

- VaR reflects risk reduction through diversification because the P&L distribution does;
- VaR reflects current market conditions again because the P&L distribution does;
- VaR has the same meaning for options, swaps, and other derivatives as it does for simpler instruments like stocks and bonds;
- A single VaR can be computed for a portfolio of diverse instruments.

This is not to suggest that VaR is the last word on risk measurement; but by explicitly introducing consideration of probabilities it appears to be a step in the right direction.

J.P. Morgan's RiskMetrics

A milestone in the adoption of VaR as a risk measurement tool was J.P. Morgan's decision in May 1995 to make its proprietary RiskMetrics system freely available through the World Wide Web. Indeed, until recently the RiskMetrics VaR Calculator was the first item on the J.P. Morgan home page, reflecting the importance the firm attached to it.[5]

In its original form, RiskMetrics consisted of

- a simple methodology for calculating VaR;
- extensive downloadable datasets (mainly estimated standard deviations and correlations for many assets);
- a 280-page downloadable technical document explaining the methodology and the datasets;
- an on-line VaR calculator that applied the methodology to portfolios of cash positions in foreign currencies.

[5] Subsequently, Morgan entered a joint venture with Reuters to further develop risk management systems. The newly formed RiskMetrics Group has its own web page at www.riskmetrics.com.

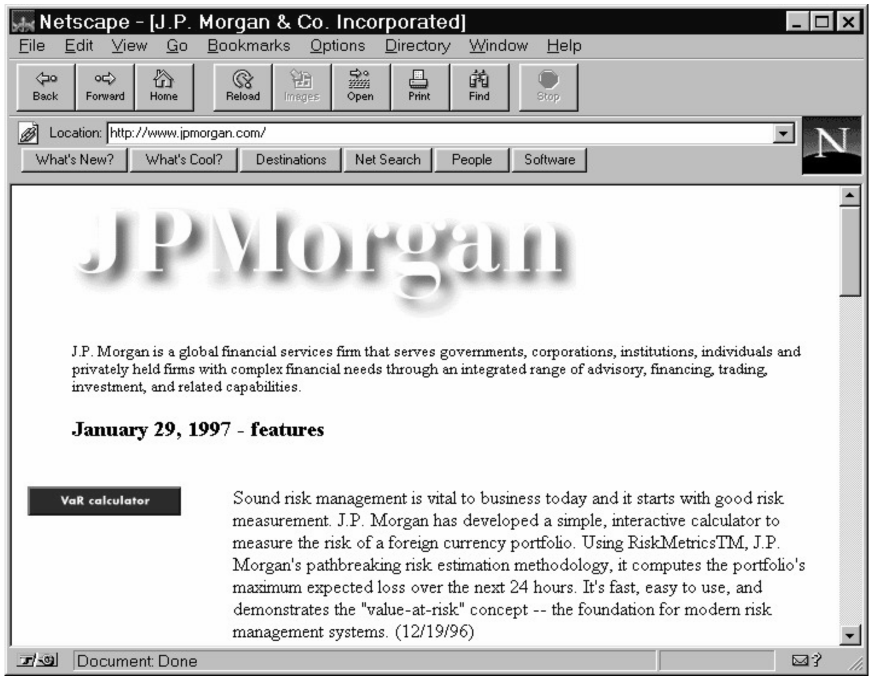

Figure 2: J.P. Morgan made its VaR Calculator the lead item on its home page

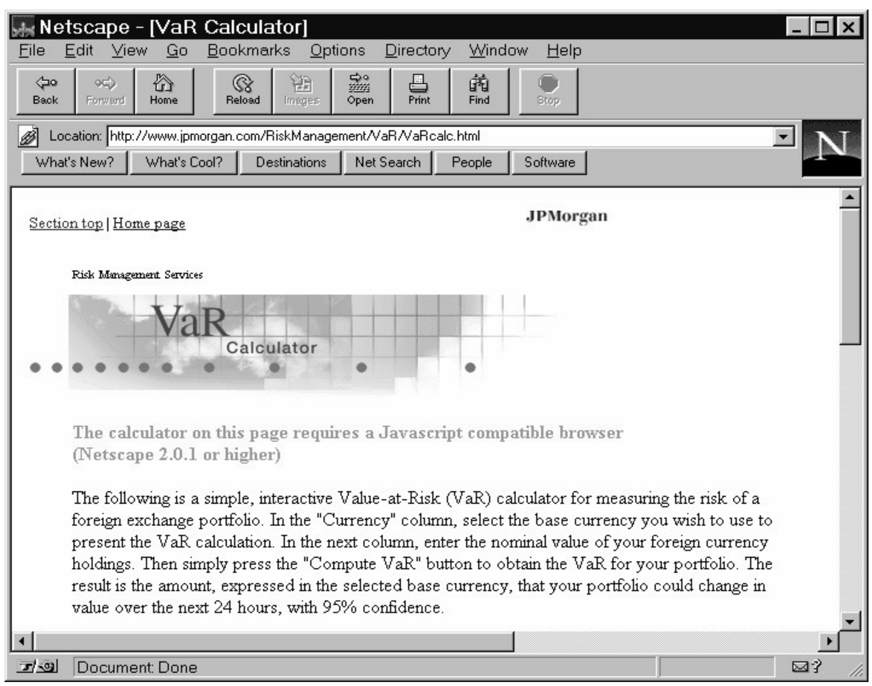

Figure 3: Explanation of the VaR Calculator

Soon after J.P. Morgan made its system available, many third-party vendors developed risk management systems using the RiskMetrics data and enhancing its methodology.

The RiskMetrics methodology entails several major statistical and financial approximations that make it easy to use but limit its accuracy. Among the most important assumptions is that all asset returns are normally distributed. This may be a reasonable (if imperfect) approximation for basic assets like stocks and currencies, but it is altogether unpalatable for, e.g., options, which typically have highly skewed return distributions.[6]

An important consequence of the assumption of normality is that *VaR can be calculated as a multiple of the standard deviation of the P&L distribution.* The appropriate multiple depends on the desired level of VaR (e.g., 95% or 99%), a higher level requiring a larger multiplier. The RiskMetrics system is essentially a way of calculating the standard deviation for a portfolio (using historical data) and then scaling it by an appropriate multiplier.

This idea is illustrated in Figure 4, an example taken from the RiskMetrics Technical Document. The "portfolio" in this example consists of just one asset (a position in German marks). The VaR is a multiple of the standard deviation of the mark/dollar exchange rate. This example uses a multiplier of 1.65 for a 95% VaR. Other multipliers are illustrated in Figure 5, taken from *Risk Management: A Practical Guide*, published by the RiskMetrics Group.

[6] In fairness, J.P. Morgan stated from the outset that RiskMetrics was not applicable to portfolios with options.

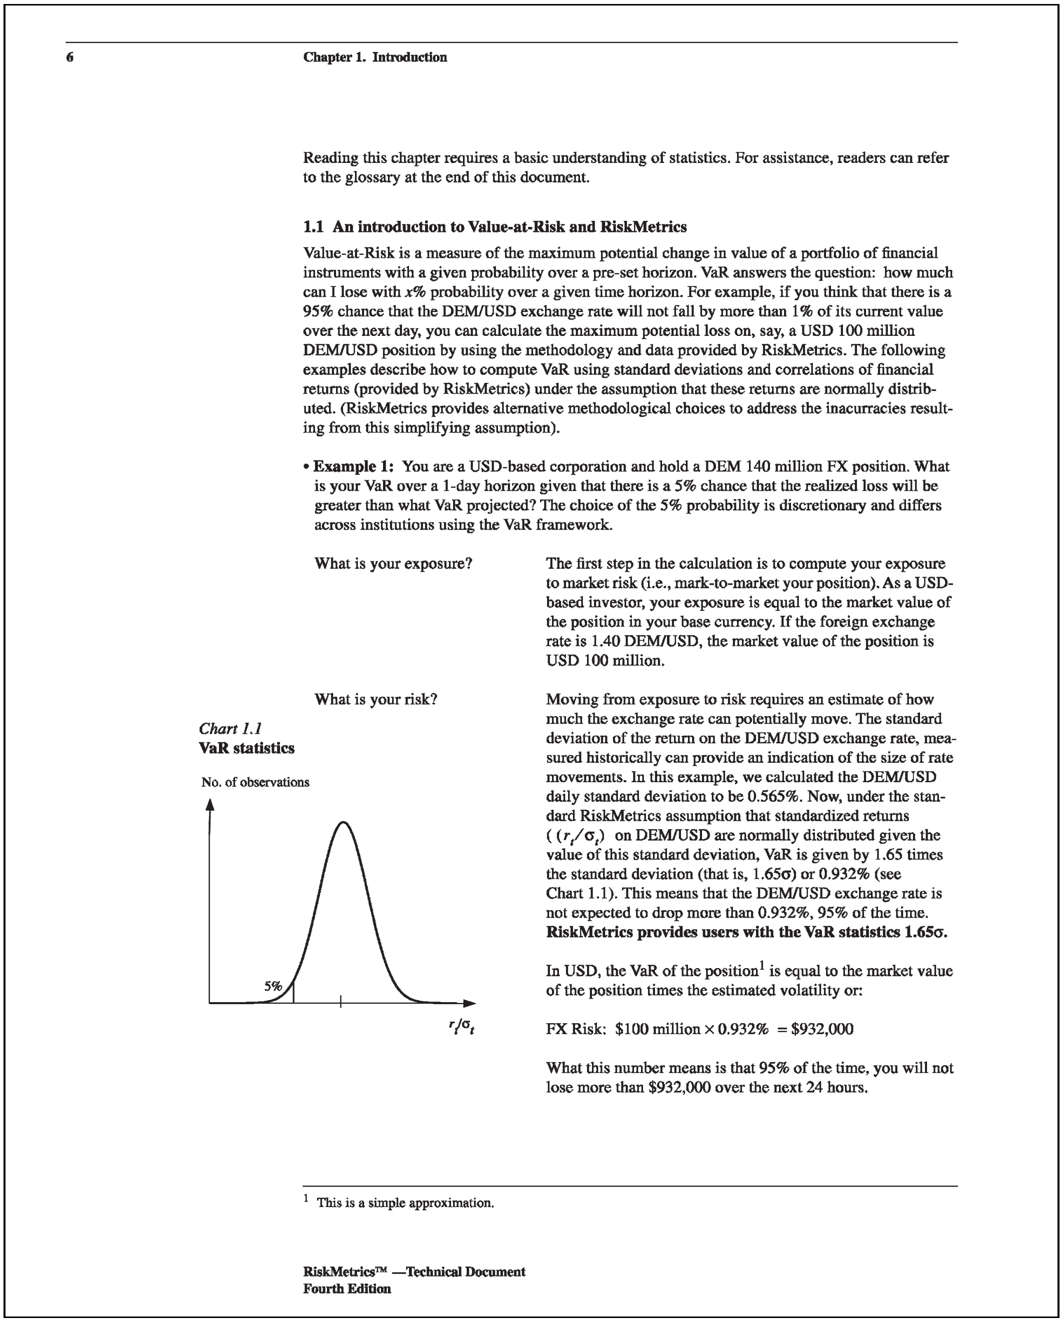

6 Chapter 1. Introduction

Reading this chapter requires a basic understanding of statistics. For assistance, readers can refer to the glossary at the end of this document.

1.1 An introduction to Value-at-Risk and RiskMetrics

Value-at-Risk is a measure of the maximum potential change in value of a portfolio of financial instruments with a given probability over a pre-set horizon. VaR answers the question: how much can I lose with *x%* probability over a given time horizon. For example, if you think that there is a 95% chance that the DEM/USD exchange rate will not fall by more than 1% of its current value over the next day, you can calculate the maximum potential loss on, say, a USD 100 million DEM/USD position by using the methodology and data provided by RiskMetrics. The following examples describe how to compute VaR using standard deviations and correlations of financial returns (provided by RiskMetrics) under the assumption that these returns are normally distributed. (RiskMetrics provides alternative methodological choices to address the inacurracies resulting from this simplifying assumption).

- **Example 1:** You are a USD-based corporation and hold a DEM 140 million FX position. What is your VaR over a 1-day horizon given that there is a 5% chance that the realized loss will be greater than what VaR projected? The choice of the 5% probability is discretionary and differs across institutions using the VaR framework.

What is your exposure?

The first step in the calculation is to compute your exposure to market risk (i.e., mark-to-market your position). As a USD-based investor, your exposure is equal to the market value of the position in your base currency. If the foreign exchange rate is 1.40 DEM/USD, the market value of the position is USD 100 million.

What is your risk?

Moving from exposure to risk requires an estimate of how much the exchange rate can potentially move. The standard deviation of the return on the DEM/USD exchange rate, measured historically can provide an indication of the size of rate movements. In this example, we calculated the DEM/USD daily standard deviation to be 0.565%. Now, under the standard RiskMetrics assumption that standardized returns $((r_t/\sigma_t)$ on DEM/USD are normally distributed given the value of this standard deviation, VaR is given by 1.65 times the standard deviation (that is, 1.65σ) or 0.932% (see Chart 1.1). This means that the DEM/USD exchange rate is not expected to drop more than 0.932%, 95% of the time. **RiskMetrics provides users with the VaR statistics 1.65σ.**

Chart 1.1
VaR statistics

No. of observations
5%
r_t/σ_t

In USD, the VaR of the position[1] is equal to the market value of the position times the estimated volatility or:

FX Risk: $100 million × 0.932% = $932,000

What this number means is that 95% of the time, you will not lose more than $932,000 over the next 24 hours.

[1] This is a simple approximation.

RiskMetrics™ —Technical Document
Fourth Edition

Figure 4: Example from the *RiskMetrics Technical Document*

1.4 Confidence level scaling factors

Standard deviations can be used to estimate lower-tail probabilities of loss when the parametric approach to measuring risk is used. Lower-tail probability of loss refers to the chance of loss exceeding a specified amount.

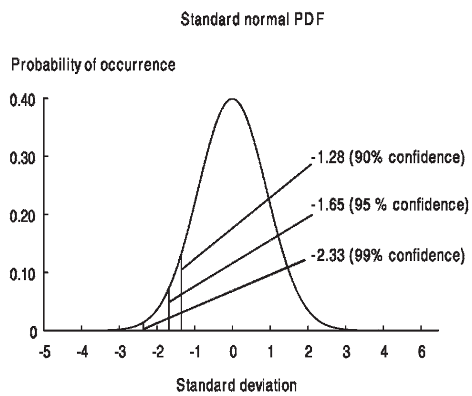

Because returns tend to cluster around the mean, larger standard deviation moves have a lower probability of occurring. To arrive at the tail probability of loss levels and implied VaR confidence levels, we use standard deviations (confidence level scaling factors). This chart shows three confidence level scaling factors and their associated tail probability of loss levels.

RISK MANAGEMENT: A Practical Guide **RiskMetrics Group**

Figure 5: Standard deviation multipliers for different levels of VaR. From *Risk Management: A Practical Guide*, RiskMetrics Group.