

PRACTICAL RISK MANAGEMENT FOR EQUITY PORTFOLIO MANAGERS

BY G. C. HEYWOOD, J. R. MARSLAND, AND G. M. MORRISON

[Presented to the Institute of Actuaries, 28 April 2003]

ABSTRACT

The paper highlights the role of risk budgeting — how risk is ‘spent’ — in the investment management process and some of the practical issues encountered. Risk budgeting has received a great deal of interest from the investment management community recently, but no clear consensus has emerged on how it should be implemented. In this paper we outline a pragmatic risk budgeting method that can be applied at the portfolio level, and show that it can produce superior results when used in conjunction with cluster analysis techniques. There are practical implications for chief investment officers and chief executive officers on how they allocate human resources and capital in the investment management process.

A statistical factor model for stock returns is used to build a risk model of the market that separates the factor components (representing the market, investment themes and styles) and the stock specific component. Then cluster analysis techniques provide a visualisation of the changing risk structure of the market. Natural groupings of stocks emerge within the market often different to the classical industrial classification systems widely used today. These natural groupings clearly change over time reflecting the changing nature of equity markets, e.g. these techniques show very clearly the emergence of the telecommunications, media, technology phenomenon in the late 1990s and its subsequent demise in early 2001.

Using the framework of a statistical factor model, risk budgets can be aggregated or dis-aggregated. Aggregation can be to country, sector or any other group. Dis-aggregation will be to common factors (e.g. the market, growth, value and other styles) and stock specific factors, derived from a multi-factor model.

KEYWORDS

Risk; Budgeting; Hedge Funds; Factor Analysis; Dendrograms; Cluster Analysis

CONTACT ADDRESS

G.M. Morrison, M.A., F.I.A., Commerzbank Securities, 60 Gracechurch Street, London EC3V 0HR, U.K. Tel: +44 (0)20-7653-7642; Fax: +44 (0)20-7645-7442;
E-mail: g.m.morrison.70@cantab.net

1. INTRODUCTION

1.1 *Overview*

1.1.1 In recent years financial markets have been subject to a great deal of change. Some examples of these changes include the following:

2 *Practical Risk Management for Equity Portfolio Managers*

- (1) We have witnessed a high, and perhaps unprecedented, level of uncertainty in investment markets.
- (2) There have been changes to society's attitude towards the provision of saving, provision for retirement and ill-health, and so on. Owners of assets require ever greater levels of accountability from their advisers. The United Kingdom Government has sponsored various reviews which are likely to have far reaching consequences.
- (3) Changes in legislation and professional standards have clearly had very material effects on the way in which financial markets work.
- (4) We have witnessed several 'bubbles' — the rise (and subsequent fall) of the technology, media and telecommunications sector and many corporate excesses, some of which have not been matched in recent memory, or perhaps ever.
- (5) There has probably never been a broader variety and choice of savings, investment and speculation opportunities available to the sophisticated and unsophisticated investor alike.

1.1.2 Market participants have reacted to these changing times in different ways, as they struggle to adapt to the new circumstances in which they find themselves. The challenges are made worse, as markets have had to make their adjustments against a back-drop of difficult and volatile markets conditions.

1.1.3 Against this background, there is a very clear demand for financial decision making to become more structured, more disciplined, justifiable, repeatable, and so on. This represents a great opportunity for the actuarial profession, whose financial modelling skills are an integral part of the examination process.

1.2 *The Aims of this Paper*

1.2.1 Against this complex background, this paper primarily focuses on the efficient management of equity portfolios, and aims to provide a practical rationale to help portfolio managers answer a number of questions:

- (1) How are the risks in any particular equity market changing over time?
- (2) How can one construct an equity portfolio more efficiently in a systematic, effective way, particularly with regard to the risk/reward trade-off?
- (3) How can one 'budget one's risk' more effectively in a practical sense?

1.3 *The Structure of this Paper*

1.3.1 Risk management of investment portfolios has never had as much attention as it has currently, yet the discipline is evolving and changing. In Section 2 we cover the topic of risk measurement, risk management and the changing paradigm of the effects of fully integrating risk management into a fund management investment process.

1.3.2 The building of efficient risk models is the fundamental building block of the entire structure. In Section 3 we categorise risk models into six broad categories: variance/co-variance, historical, factor, value-at-risk, statistical and stochastic models. We review the different types of risk model in turn.

1.3.3 The building of a risk model is a relatively technical operation, but it is essential for practical application for portfolio managers, even though it is of little interest to them in its own right. In Section 4 we overlay the application of cluster analysis onto the risk model in an innovative way to show the risk structure of the market. Repeating this process at various points in time shows how the risk structure of a market changes over time. This is reviewed in Section 5.

1.3.4 To an investment manager ‘risk’ is an important and precious commodity, and needs to be spent prudently. Whilst a number of papers have been written on this important topic, there seems to be a dearth in the current literature on practical applications to help fund managers in this task. In Section 6 we review the literature and provide a series of practical examples of how a statistical factor model combined with cluster analysis techniques can form a useful and practical toolkit.

1.3.5 Our conclusions are covered in Section 7.

2. THE CHANGING PARADIGM OF PORTFOLIO MANAGEMENT

2.1 *Background to Risk Management*

2.1.1 Managing a portfolio is essentially a process of balancing expected risks and expected returns, bearing in mind any restrictions and constraints that there might be. For example, these constraints may be placed on the portfolio manager by the client, a regulator, or may be effectively self-imposed by the fund manager for professionally prudent reasons.

2.1.2 Arbitrary stock restrictions were historically a way of managing risk. Popular ways of managing risk in the past have utilised *ad hoc* portfolio construction rules. A restriction of a maximum of 10% in any single stock within a portfolio is a good example of this. Other examples include minimum and maximum sector and country weights, often relative to a client-specified benchmark; restrictions on large vs. small capitalisation exposures, and so on.

2.1.3 Indeed, it could be argued that the widely practised idea of sector neutrality within a portfolio is just another arbitrary restriction. However, as we will show later, ensuring the money neutrality of holdings within a portfolio is a wholly different thing to true risk neutrality. For example, a fund might have the same exposure to the telecommunications sector as in the fund’s benchmark, but if you own the more volatile stocks in the sector and do not own the less volatile telecommunications stocks, the likelihood is

that you will have a risk ‘overweight’ position in the telecommunications sector as a whole.

2.1.4 Furthermore, there has been no widespread recognition of the idea that portfolio construction should be a distinct discipline within a fund management house; the problems and challenges of portfolio construction and risk management are every bit as much a ‘science’ as they are an ‘art’. Frequently, the same professionals who are involved in asset allocation or stock selection are often involved in portfolio construction, without accepting that the skill sets required may be very different. It is our view that fund managers will increasingly make this distinction in the future. More recently, the legal and regulatory systems, combined with a more sophisticated end client base, have started the process of making risk an explicit constraint on the portfolio manager. Consultants and regulators are also becoming more and more interested.

2.1.5 Modern portfolio and risk analysis systems give the fund manager the tools to manage risk and return interactively, allowing them to both comply with regulations and client restrictions and to best manage return under these constraints. In our view, it is essential for fund managers to have such systems in the face of increasing competition and client accountability. This is distinctly different from simply monitoring risk, say once a month in arrears, to comply with the minimum standards of due diligence expected by the client.

2.2 *Definitions of Risk*

2.2.1 We believe that it is important to distinguish between three definitions of risk; these differences are more than semantics.

2.2.2 *Risk monitoring* is most frequently — but not always — observed in arrears. Typically, it will answer some, or all, of the following questions:

- What level of risk has been incurred?
- What were the sources of that risk?
- Has any unexpected risk been incurred?

2.2.3 Clearly, the main drawback to this definition of risk monitoring is that risk is observed in arrears; by which time the risk has been incurred, it cannot be managed, and there is nothing that the portfolio/risk manager can do about it.

2.2.4 *Risk measurement* is the act of measuring the level of risk, and is unique to the portfolio, its benchmark and the risk model used. Clearly, the objective is to measure this risk as accurately as possible.

2.2.5 *Risk management* is practised in real time. Typically, it is a separate process to risk measurement, and the ‘risk views’ are not unique. Risk management aims to answer the following questions:

- What is the level of expected risk that is being incurred?
- Does the asset portfolio capture the desired risks?

- Is any unwanted (unexpected) risk being incurred?
- What is the impact on expected risk if the portfolio is changed?
- Is the expected risk/return payoff acceptable and efficient?

2.2.6 Integrating risk management into the portfolio manager's daily portfolio construction process is both a significant improvement on arbitrary stock/sector/country restrictions and a step improvement over the occasional measurement of risk by an external team. Empowering portfolio managers with the tools to manage risk should allow them to add value in the form of better managing the risk/return characteristics of their portfolios.

2.2.7 Figure 1 starkly shows the differences between the processes underlying risk measurement and risk management.

2.2.8 The representation of a risk measurement system in Figure 1 can be summarised as follows:

- It is relatively simple and can encompass any type of return forecasting, whether explicitly quantitative or more subjective and traditional. Often, we accept, portfolio managers do not have explicit expectations of return, but merely some sort of preference ordering.
- Portfolio construction drives the risk checking process; there is no interaction and iteration between the two processes.
- Trading is not a part of the portfolio construction process.

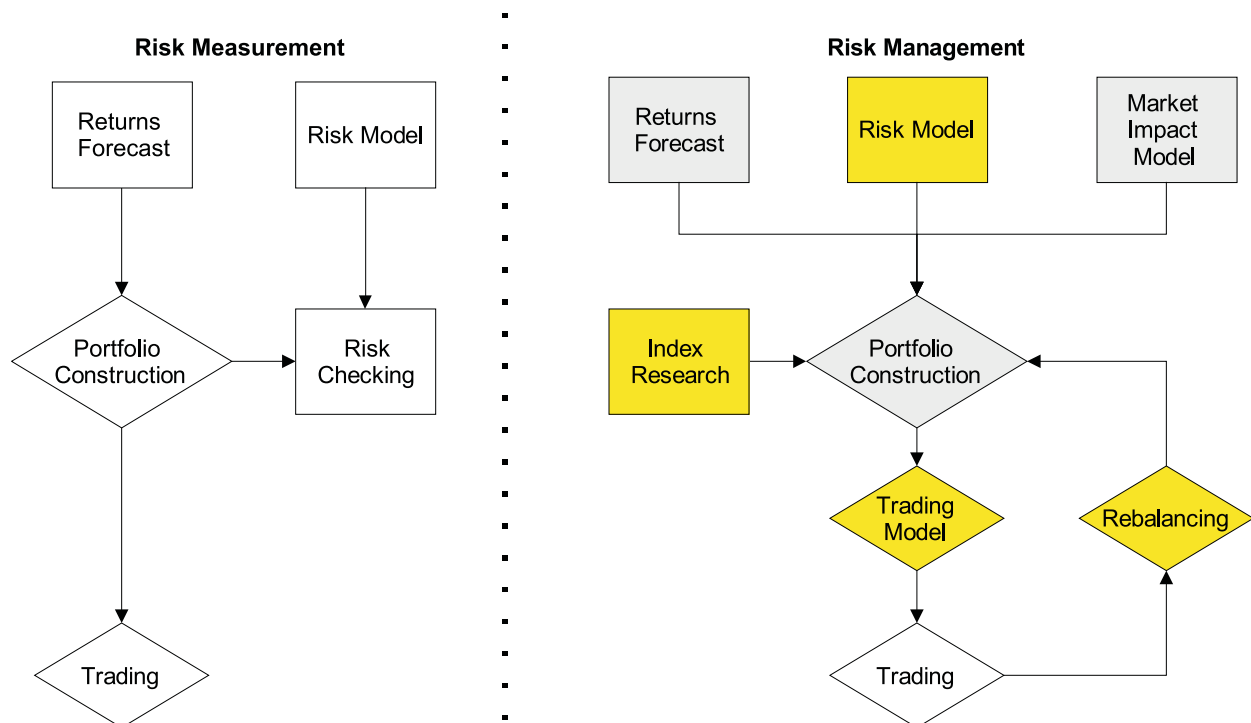


Figure 1. Risk — the differences between risk measurement and risk management

2.2.9 The processes underlying risk management are altogether more involved and complex:

- Portfolio construction is at the heart of the process, with direct input from risk models, market impact models and index research.
- Trading models assist the process of implementation, and naturally re-cycle into the portfolio construction process.
- A market impact model is an independent and explicit input into the management process — the cost of trading has to be explicitly incorporated in the risk/return trade-off aspects of the portfolio construction process. We have not expanded on this important topic in this paper.

2.3 *Why is Risk Management so Vital?*

2.3.1 Many fund managers have not practised risk management. Risk management for a financial enterprise requires both the aggregation of positions across asset classes and the understanding of risks inherent in those positions. This is by no means a trivial or simple task, even for modern organisations with access to sophisticated information technology, which may control many hundreds of thousands of positions invested in a variety of instruments traded across different time zones. In order to obtain a solution to this difficult management challenge, simplifications have typically been made, both at the level of aggregation and of risk analysis. While these simplifications may be unobjectionable for some purposes, certain applications demand a more sophisticated approach.

2.3.2 The clear implication is that the technology applicable to a single portfolio, typically based on a multiple factor model for security returns, must be implemented enterprise wide.

2.3.3 The asset management industry has become increasingly complex over recent years. Organisations that may have focused, a few years ago, on delivering one or two similar products constructed in one location to a single, homogenous group of clients, have evolved into true multinational enterprises.

2.3.4 Portfolio ‘manufacturing’ may take place in a number of locations (the equivalent for a multinational is a factory) around the world, and in each location different styles and variations of the product line may be developed.

2.3.5 Particularly in the large fund managers, manufacturing competencies are kept distinct from the skills required to develop an efficient distribution capability.

2.3.6 Portfolio distribution is likely also to be a multi-location activity, and in each location a variety of different channels may be employed to reach customers of interest to the organisation.

2.3.7 In many regards, the modern fund manager is as organisationally complex as any large, industrial, multinational company. However, there are

four clear differences between a traditional multinational industrial company and a modern asset management firm:

- (1) The fund manager, despite the increasing use of technology, is still highly dependent on individual human contributors — the ‘assets’ of the company go up and down in the lift each day. Traditionally, investment has always had a heavy reliance on ‘investment flair’ in a way that has little parallel in an industrial company.
- (2) There are a number of differences in the level of regulatory oversight of the investment product. At the ‘factory’ — where prudence and fiduciary standards are key operating constraints — the level of regulation of the ‘product’ is relatively low. (This contrasts to the industrial company where freedom to manufacture might be constrained by patent and safety laws.) However, as far as distribution is concerned, the regulation is far more restrictive, and securities laws govern the transparency of the sales process.
- (3) This regulatory environment, at least historically, imposes a structure, which has forced a significant human intervention. In many instances it is as if each product and sales effort is individually ‘hand made’ for the ultimate client or prospect. Fund managers have tried to mitigate this development by increasing, where it is possible and practical, the homogeneity of portfolios, e.g. by establishing commingled vehicles.
- (4) With so many ‘moving parts’, the management and control problem of a modern fund manager is enormously complex. He or she has to co-ordinate intelligent, motivated individuals, who, in many cases, represent the ‘value’ of the organisation, to perform an intricate task efficiently and to retain some scope for personal challenge and reward. Meanwhile, the organisation needs to overlay a structure for achieving stability and growth to ensure product quality, at the same time delivering a return on capital for the shareholders.

2.4 The Role of Portfolio Construction and Risk Analysis

2.4.1 In addition, the portfolio’s overall risk and the portfolio’s likely deviation from its benchmark (tracking error or residual risk), among other risk statistics, are important, not only from a quality control point of view, but also to satisfy the regulatory requirements.

2.4.2 There are many good examples of how an incomplete knowledge of the aggregate risks can lead to inefficiencies in the risk budgeting process, or worse. There are three broad classes of problems which arise from a failure to have proper risk management systems. These are:

- (1) a potential compounding of unintended bets and uncompensated exposures to risk;
- (2) unwarranted over-diversification or the reverse, concentration; and
- (3) a shorter window misallocation of resources.

2.4.3 *Unintended bets*

- (1) In an organisation managing a number of different portfolios, a single source of risk can impact the portfolios underlying many of the products offered, with a huge potential impact on the overall assets. This is true both in a domestic as well as in a global context.
- (2) For example, recent events in global financial markets have clearly shown that volatility ‘flows’ around the world, and that events in a single market can have important effects in a global context. Recent specific examples of these phenomena relate to global managers with strong cultural tendencies to a focused style, such as value or growth managers under-performing throughout their asset base, depending upon which style is in favour.
- (3) Without aggregating portfolio holdings and critically examining their sensitivity to risk factors, no chief investment officer (or chief executive officer) can be certain of the bets that are being made. Moreover, the decision makers cannot be reassured that the bets, in aggregate, are likely to be compensated by returns sufficient to justify their risk.

2.4.4 *Over/under-diversification*

A lack of co-ordination between multiple managers can lead to three potentially undesirable outcomes:

- A series of unintentional small bets can compound into an unwarranted large bet, leading to an unjustified over-concentration in the aggregate portfolio. In contrast, without co-ordination, the portfolio managers may intentionally take a similar bet (for example, towards value stocks world-wide), leading to an aggregate bet that is far too large and an over-concentration in the aggregate portfolio.
- The opposite, over-diversification, is also a danger. Here conservatism tends to eliminate all bets, with the aggregate portfolio approximating an index fund without the possibility of superior performance.
- For a fund of funds or a plan sponsor portfolio, active management fees go un-rewarded, while, for an asset management firm, the possibility of eye-catching performance is eliminated, together with any justification for management fees.

Clearly, similar problems can occur within individual portfolios and their sub-components.

2.4.5 *Misallocation of resources*

- (1) Modern financial theory is centred on the goal of maximising a risk/reward trade-off. This objective applies just as much to the enterprise as a whole as it does to an individual portfolio, where quantitative managers, in particular, have long used analytical tools to ensure that an asset’s contribution to portfolio risk is commensurate with its contribution to expected portfolio return.

- (2) We believe that a similar paradigm applies to a fund of funds (in terms of the component portfolios), a plan sponsor (in terms of the managed sub-portfolios), and an entire asset manager (in terms of the component products). Just as the success of an individual portfolio manager is tied to the performance of the portfolio and its risk/reward trade-off, so the long-term success of the asset management firm is tied to the aggregate performance of all the products offered and their aggregate risk/reward trade-off.

2.5 What is Risk?

2.5.1 Risk is often a misunderstood concept, with no clear conceptual or quantitative acceptance of what it means, let alone how it could be measured or even managed.

2.5.2 Risk is clearly a multi-faceted concept, that means:

- different things to different managers;
- different things to the same manager at different points in time;
- institutions will have different risk tolerances;
- institutions will price risk at various levels; and
- risk can be interpreted in different ways.

2.5.3 It is likely to depend of the objectives of the risk measurer, his financial position, as well as the time horizon over which he wants to measure and manage his risk.

2.5.4 A number of examples looking at risk from different perspectives (looking at the same problem through different ‘spectacles’) will clarify the point. In order to bring this idea to life, we look at the nature of risk in four generic investment institutions in turn:

- an insurance company;
- a hedge fund;
- a mutual fund; and
- a U.K. pension fund — as perceived from the different interested parties involved.

Clearly, a similar thought process and analysis can be extended to other types of institutions and market participants.

2.5.5 Insurance company

- (1) Consider an insurance company and the various interested parties: policyholders, the company and its shareholders, and insurance industry regulators.
- (2) The policyholders’ view of risk centres on the insurance company’s ability to meet the claims in the event of the insured event occurring. For policies with a with-profits element, policyholders will have reasonable expectations with regards to investment gains.

- (3) The company and its shareholders have to consider the broader asset/liability risk environment. They will be looking to maximise their returns, given some definition of risk. Typically, they have a long-term horizon.
- (4) The regulators will take a rather different view. They are primarily concerned with the financial strength of the insurer, and will want to see in it as high a solvency level as possible. There will be a tendency to be cautious and prudent, and they will naturally tend to constrain the level of risk that can be borne. The regulator will typically be very keen to segregate the policyholders' funds from those of the shareholders, and in many markets this is a legal requirement.

2.5.6 *Hedge fund*

- (1) In sharp contrast to insurance companies, hedge funds are established in such a way that they are subject to relatively low levels of regulation, thus giving them large amounts of investment freedom. Typically they have the fewest investment constraints of any category of investors; by and large they are free to take on whatever level of risk they feel comfortable with, including, importantly, leverage. Leverage is employed by hedge funds in different degrees and for different purposes, among them increasing either the size or the number of positions in the fund's portfolio, amplifying the small residual returns generated by spread trades and offsetting the fund's directional market exposure.
- (2) Hedge fund managers can differ substantially on how they implement their strategy and are free of the constraints of being measured relative to a benchmark index. As a result, the correlation of returns between different managers and different strategies is frequently low and stable. Typically, fund-of-fund strategies are structured to exploit these features.
- (3) Most hedge funds define risk as a loss of principal as opposed to tracking error relative to a benchmark (Parker, 2002).
- (4) The time horizon of this risk is typically much shorter than for many other types of portfolios; it is not unusual to measure and manage this 'draw-down' of capital on a daily basis.
- (5) Rate of return is normally measured relative to cash and a 'peer group' of hedge funds with a similar style.

2.5.7 *Mutual fund*

- (1) Similar to insurance companies (and unlike hedge funds), mutual funds are aimed towards the retail client, and consequently are relatively highly regulated. They operate in a highly competitive environment, where assets have a tendency to move towards the 'fashionable' fund manager or 'hot' investment sector.
- (2) The rate of return is of crucial importance. However, it tends to be viewed as relative to the competitive 'peer group' rather than either a benchmark or a cash deposit rate.

- (3) The risks are usually seen as not performing as well as the ‘peer group’ — perhaps more a manifestation of a business, rather than an investment, risk.
- (4) The time horizon is relatively short, but, typically, it falls between the short-term focus of the hedge fund and the longer-term perspective of either a pension fund or an insurance company.

2.5.8 *Pension funds*

Consider a U.K. pension with defined benefits, where the benefits are linked to salary close to retirement; the assets are managed externally; there is an independent group of trustees, some of whom ‘represent’ the employees. The external managers have been set investment objectives, guidelines and constraints by the trustees, acting on the advice of an independent investment adviser. These objectives may be expressed in the form of rates of return and acceptable levels of ‘risk’, perhaps expressed in the form of tracking error, asset allocation, concentration of investment, and so on. It would not be surprising if the interested parties looked at the risk in the following (different) ways:

- (1) *The trustees.* The pension fund trustees have only the interests of the scheme members in mind. They view risk solely from the perspective of the fund’s ability to pay future benefits and any pension increases or benefit improvements that might be made.
- (2) *The company.* The finance director is likely to be concerned about the cost of the pension fund, the potential level of volatility in the pension expense and the effect that this may have on the profit & loss statements, the balance sheet, etc. This concern will extend to how the company is perceived in the market place by investors, analysts, credit-rating agencies and bankers. On the other hand, pension provision is seen, quite rightly, as an important part of employees’ total remuneration packages.
- (3) The company is likely to be more short-term focused, with ‘risk’ closely inter-twined with the potential increase in contributions and the effect of volatility on the corporate accounts.
- (4) The finance director will become increasingly interested in risk modelling the combined company/pension fund/wider employee benefits structure as a single entity — particularly in the light of FRS 17, the new minimum funding requirement and the changing pattern of pension provision.
- (5) *The fund manager.* The fund manager will be driven to look at risk in yet further different ways.
- (6) He will obviously view the chance of not meeting the targets set out in the investment management agreement as a risk. It is possible that he will have a different perception of the risk of not meeting those targets as being different to exceeding them — his risk may be asymmetric.

- (7) He also has a business risk relative to his ability to attract new funds under management; how well has he performed (in the broadest sense) relative to his competition.
- (8) The investment risk that the manager is prepared to bear will depend on his views of the markets. When he has a clearer opinion on the risk/reward benefits offered by the markets he may wish to 'turn up' the risk level, and vice versa when the investment climate is more uncertain. Furthermore, the nature of the risks perceived by an individual investment manager are likely to vary from time to time, and at any point in time different fund managers will have a different perception of the risks within the market. In a large pension fund (or a fund-of-funds) it is likely that the views of many portfolio managers with different views will be aggregated within one fund, but across many geographic regions and asset classes.

3. REVIEW OF RISK MODELS

3.1 *What is a Risk Model?*

3.1.1 A risk model is the key ingredient that allows a portfolio constructor to put his expected returns for different stocks into context.

3.1.2 Typically, the portfolio manager has to consider a number of things when efficiently constructing a portfolio:

- the sources of return for each stock;
- the sources of risk; and
- the concentration of the portfolio and the diversification of the sources of risk and return.

3.1.3 Risk is often not as intuitive as return, because it is multi-dimensional. Risk models seek to simplify the problem by allowing the portfolio manager to make more sensible use of the available information. Again, we confine our attention to the risk models that can be used by portfolio managers.

3.2 *Criteria for Choosing a Risk Model*

3.2.1 There are multitudes of different ways in which risk models can be built. In our view, there is no 'correct' methodology that can be applied in all circumstances. However, some models and methodologies are better than others. There are four criteria that can be applied in the assessment of the quality of risk models:

- explanation of risk;
- objectivity;
- interpretability; and
- forecasting of risk.

3.2.2 *Explanation of risk* is the ability to break risk down into a lower number of more or less independent dimensions. Historical, Monte Carlo and variance/covariance techniques do not attempt to simplify risk, only to measure it. Factor models break risk down into common factors and stock specific components. Common factors should ideally be independent, i.e. contain uncorrelated information. A measure of success of these models is how much of the total risk of each stock can be explained by the factor component. Typically, macro-economic factors explain the least proportion of risk, and statistical factor models, by definition, have maximum explanatory power for a given number of factors.

3.2.3 *Objectivity*. It is necessary to make a number of assumptions when building risk models. Objectivity implies a lack of assumptions about what drives the differential performance of stocks. A macro-economic factor model relates stock returns through their sensitivity to prevailing economic forces. Similarly, fundamental models relate the characteristics of a particular company to corresponding common risk factors; for example, relating the market capitalisation of a company to a market size factor return. Each of these factor models makes assumptions about the precise identity of the common factors driving returns. Obviously, it makes no sense to test theories about what drives stock returns in models that have made assumptions about these processes beforehand. Statistical factor models, unlike other types of risk model, make no such a priori assumptions about the precise sources of risk.

3.2.4 *Interpretability* is often the flip side of objectivity. Macro-economic and fundamental factor models have the advantage of relating real world risk factors to stock price returns. For example, a macro-economic model might specify how each would respond to an unexpected change in the rate of consumer price inflation.

3.2.5 *Forecasting of risk* for historical, variance/co-variance, macro-economic and fundamental risk models implicitly assumes that the historical risk structure of the market will, on average, continue in the future. We know this assumption to be invalid and, in certain market environments, extreme events, or even slow trends, can introduce substantial errors. Monte Carlo techniques can explicitly make forecasts for the future structure of the market, but there are large subjective elements in the distribution assumptions. Statistical factor models ideally lend themselves to forecasting techniques, and the forecasts have the further benefit of being objectively driven by the data.

3.2.6 Mixed factor models seek to combine the advantages of each of the three main factor modelling techniques, namely macro-economic, fundamental and statistical.

3.3 *Types of Risk Model*

3.3.1 There are a number of differences between the underlying

approaches to constructing risk models. We classify six different types of risk model, and briefly consider each in turn:

- variance/co-variance methods;
- historical models;
- factor models;
- value-at-risk models;
- statistical models; and
- Monte Carlo techniques.

3.3.2 *Variance/co-variance methods*

- (1) Variance/co-variance methods are based on the work done in the 1960s by Markowitz (1959) and Sharpe (1963). These models formed the basis of modern portfolio theory.
- (2) Given a number of assumptions, that more modern techniques are able to relax, the correlation between assets can be allowed for in measuring the overall riskiness of a portfolio.
- (3) The main problem is that the model is purely descriptive, and does not allow the sources of risk to be analysed, and so renders them useless for modern risk management. Furthermore, the assumptions underlying the model are too restrictive for more modern assets like derivatives.
- (4) Many other types of risk model are based upon this fundamental approach. For example, both pre-defined and statistical factor models typically decompose the historical co-variance matrix in terms of a parsimonious set of factors.

3.3.3 *Historical models*

These models typically use achieved returns on portfolios to estimate the risk that has been incurred in the management of a portfolio. Thus they are concerned with risk measurement rather than risk management.

3.3.4 *Factor models*

- (1) Factor models seek to explain risk by building on the variance/co-variance approach and adding explanatory structure in the form of different factors (see Ross, 1976). There is great choice of explanatory variables, but they fall into two broad categories. The factors are typically either macro-economic or fundamental.
- (2) Macro-economic factors essentially try to model the sensitivity of equities and other assets as a function of economic factors. The most common factors are usually:
 - interest rates (short-term, long-term, shape of the yield curve);
 - currencies;
 - inflation (consumer prices, producer prices, unit labour costs);
 - commodity prices (oil, gold, indices); and

- output (gross domestic product, industrial production, retail sales, survey data)
- (3) Fundamental factors are generally based upon data derived from corporate accounts, and are felt by the investment community to be important factors that drive equity prices from time to time.
- (4) Fundamental factor models express the riskiness of assets as a function of various styles and indices. The most common factors are usually:
 - value vs. growth (price/earnings ratio, price-to-book, yield);
 - the size (log market capitalisation, 'blue-chip' effect);
 - momentum/success (index out-performance, moving averages);
 - forecasts/surprises (I/B/E/S expectations, earnings revisions);
 - and
 - the country or economic/industry sector effects.
- (5) Despite its undoubted popularity, this type of model is fraught with a number of serious problems. The models intrinsically lack flexibility; they do not respond well to changes in market conditions or to new variables that may drive prices. In most cases the factors simply do not match up to those that are used by the portfolio managers. There are a limited numbers of factors; different factors would require a completely new re-estimation of the model that often renders the exercise impractical. The factors are correlated, and therefore interpretation of the results, whilst it appears to be quite simple, is, in fact, extremely difficult. In the case of economic series, most economic series are highly correlated, and one runs into severe problems when including many factors. Frequently, meaningful data are not available on a consistent basis either across or within markets.

3.3.5 *Value-at-risk models*

- (1) This approach to modelling risk management comes from a different perspective, i.e. how much money could I lose at a given level of probability? The estimates can be based on either parametric estimates of the distribution of returns or non-parametric statistics.
- (2) Value-at-risk (VaR) has traditionally been practised by investment banks for internal risk management.
- (3) VaR models are not without their problems. Non-parametric VaR models are mainly descriptive, and do not allow the sources of risk to be analysed. The analysis is limited to simulation and scenario analyses; there are no sensitivities to factors. Finally, and perhaps most importantly, risk attribution is difficult. However, parametric VaR models are more similar to their 'cousins' in factor and statistical model categories.

3.3.6 *Statistical models*

- (1) Statistical models take an abstract approach to modelling the risk of assets. Typically, these models are based solely on market prices and dividends, and make very few assumptions as to what drives the risk in markets at any point in time.
- (2) Based upon these rates of return, statistical techniques are used to produce a set of statistical risk factors. The results are so-called ‘blind factors’, which typically have a better fit with the asset returns than with other methods that use pre-specified factors in some way.
- (3) However, there are a number of problems with this approach. These ‘blind factors’ are difficult to ascribe meaning to, and have no ‘real world’ application. Furthermore, the estimation techniques tend to require a clean and complete data set, which is difficult to achieve in practice — for example with initial public offerings, privatisations and internet stocks.

3.3.7 *Monte Carlo techniques*

- (1) Monte Carlo techniques use large numbers of randomly generated scenarios to highlight the range of possible outcomes and, therefore, risk. These types of technique are well known to the actuarial profession, as they are applied quite widely.
- (2) These methods are not different risk models as such, rather alternative ways of estimating the shape of the more complex probability distributions (e.g. non-normal, leptokurtotic and skewed distributions).

3.4 It is worth considering these differences more systematically alongside a ‘wish list’ of what a good risk model actually provides. Table 1 summarises our view of how each of the main methodologies meets the four criteria that we set out in Section 3.2 for choosing a risk model.

3.5 *A Mixed Factor Model Approach*

3.5.1 We favour a mixed hybrid approach that seeks to add interpretability to the statistical factor approach, as well as to provide an improvement on the most simple statistical models. There are a number of important features in our approach.

Table 1. Summary of differences between different types of risk model

	Factor models					
	Historical	Monte Carlo	Variance/ co-variance	Macro- economic	Fundamental	Statistical
Explanation of risk	-	-	-	✗	✓	✓
Objectivity	✓	✓	✓	✗	✗	✓
Interpretability	-	-	-	✓	✓	✗
Forecasting of risk	✗	✓	✗	✗	✗	✓

3.5.2 An important aspect of our approach is the distinction that we make between risk measurement and risk management. Inefficiencies and lack of clarity are introduced if these two aspects of risk are not analysed separately. If the model tries to do both simultaneously, then the measurement of risk is likely not to be as accurate as if the two components were kept separate.

3.5.3 Therefore, in a mixed factor model there is a two-stage approach. First, there is the measurement of the risk, and second (and quite separate), there is the interpretation of the risk so that it can form the cornerstone of risk management.

3.5.4 *Phase 1. Building the model to measure risk*

- (1) For the risk measurement phase only market prices are used. The rationale for this is that it is not possible to know, on a consistent basis, the risk factors that are driving stock markets at any point in time. The best estimate that is available is as indicated by the relative risk preferences of the market participants — this is clearly reflected in the marginal price at which these participants are prepared to transfer their preferences into their portfolios, i.e. market prices.
- (2) By using market prices, we expect to have a better measure of risk than if we had applied a pre-specified factor model.
- (3) A number of techniques are used to construct a base statistical factor model with orthogonal factors. We use maximum likelihood techniques (which we believe are more appropriate than principal component analysis for this particular application) (Dempster *et al.*, 1977). The factors represent a mathematical description of the common movements in stock returns. The factors are orthogonal to each other, i.e. they contain non-overlapping information and are uncorrelated. The resulting residual risk for each stock can be viewed as specific to that stock and unrelated to the other stocks in the model. These mathematical properties are very useful in subsequent analysis, particularly with regard to risk management and the interpretation of risk.
- (4) It is important to recognise that the maximum likelihood techniques are not new statistical techniques — there is a rich academic literature on the methods. However, it is different to the better-known technique of principal components (Shukla & Trzcinka, 1990). The main differences between the two techniques can be summarised as follows:
 - Principal components' analysis is a matrix manipulation technique, and therefore requires a complete data set of returns — it does not handle missing data well.
 - Whilst, by design, the factors derived in a principal components' analysis are orthogonal, the balancing item is only that part of the risk that is not already explained by the factors. It is not independent of the derived factors, nor are they independent of each other.

- (5) From this base model other models can be generated, which are more appropriate to the job in hand, but it should always be possible to rotate back to this common point of reference.

3.5.5 Since they are based on market prices, the modelling process is very flexible and consistent models can be created across markets — the only criterion is a set of consistent market prices. The model can be adapted to cope with assets that trade infrequently or have a short trading history — but there will be estimation error.

3.5.6 From this base model we can then start to create other models. These include:

- (1) trading or arbitrage models with different periodicities, e.g. daily;
- (2) pre-specified models, where either fundamental company or macro-economic factors can be combined with statistical factors to achieve a more meaningful and intuitive structure along with an improved model fit; and
- (3) back-testing models, using different estimation periods, to test investment strategies in the past.

3.5.7 Clearly all risk models are estimated with some sampling error. Models of different data windows and periodicities can easily be estimated and compared. However, there is always a trade-off between a more responsive model, based upon a shorter window of possibly higher frequency data, and a more stable model based upon a longer window for less sampling error.

3.5.8 *Phase 2. Managing risk*

- (1) Investment themes vary over time, and different people are interested in different themes. This is a major challenge for pre-specified factor models — whether they use fundamental or macro-economic factors — since they tend to vary slowly over time, are lethargic at capturing new influences on market returns, and are most unlikely to accurately reflect and capture a portfolio manager's investment processes or valuation disciplines. This is a big impediment to practical risk management. The mixed factor model structure allows risk to be viewed in a highly flexible fashion and to cope easily with different and transient themes.
- (2) Themes can be highly client specific or highly time specific, for example:
 - a portfolio manager's investment process might be value-based and tailored to two or three proprietary valuation measures; or
 - technology, as an investment theme, affected markets in a very different way in 1999/2000 than in previous years.

3.5.9 The statistical model provides a base from which to look at a multitude of factors without re-estimation of the model. Using a given theme

that can be specified by the portfolio manager, the process reveals the various investment exposures that have been taken in a portfolio, either in absolute terms or relative to a benchmark.

3.5.10 One might expect very different results from quite similar themes, say, for example, value growth. The textbook approach of using price-to-book value as a theme will yield very different results to, say, enterprise value/earnings before interest and tax, price-to-cash flow or earnings surprise data. The mixed model approach makes it possible to analyse all of these sub-themes.

3.5.11 We will see, in Sections 4, 5 and 6, how the features of the mixed method risk model can be used to help in practical risk management.

4. RISK STRUCTURE OF AN EQUITY MARKET

4.1 *Why is Risk Management so Vital?*

4.1.1 Much of the discussion seems to have centred on risk control, risk monitoring and very little on risk management. The differences are not just semantics — they go to the very heart of a modern investment management process.

4.1.2 Whilst risk control and risk monitoring are interesting topics in their own right, they are akin to driving a car (sometimes in ‘odd’ markets a very fast car) whilst the driver spends all his time looking in the rear view mirror — sooner or later the car will be involved in an accident. However, risk management is an integral part of a fund management process, and is practised by portfolio managers in real time, thus helping to avoid unnecessary delays in taking action on the level of portfolio risk.

4.1.3 Nevertheless, we have some sympathy with the users of risk models and associated tools today. They:

- are hardly user friendly, and are not accessible to portfolio managers in their day-to-day jobs;
- seem largely irrelevant to the risk problems facing portfolio managers;
- are often based on out-dated quantitative techniques; and
- necessitate a reasonably advanced knowledge of statistics to understand and interpret the results.

4.1.4 Against this background, it is little wonder that the generic topic of risk has been historically consigned to a back-room ‘risk controller’, whilst fund management has continued largely unaffected.

4.1.5 For risk models to be of any use whatsoever, they have to be incorporated into the risk management process. Even the most sophisticated risk models need to be able to provide:

- easily understandable tools that are available to portfolio managers to enable them to construct and monitor portfolios in real time;

- an easy way to explore the changes to the risk/reward profile of a portfolio based upon hypothetical transactions; and
- an efficient way of showing how to budget risk effectively, so that the manager can clearly identify the relative positions that he is taking in the portfolio, the subjective risks he is taking relative to the returns that he is trying to achieve, and the stocks, if any, that he has in the portfolio purely for risk management/reduction purposes.

4.1.6 It is perfectly possible to successfully play computer games — like space invaders — without having a detailed knowledge of the mathematics behind how the computer works, the program that controls the game, or how the results are visualised on the screen. So it can be with portfolio construction and risk budgeting, if an appropriate approach is adopted.

4.1.7 Both portfolio construction and risk budgeting are different aspects of the same problem. They are both concerned with constructing efficient portfolios that capture the portfolio manager's risk/reward trade-offs, subject to an acceptable level of risk (perhaps set by the 'owner' of the assets), whilst 'spending' the risk in a controlled manner. You cannot have proper portfolio construction without a rational framework of risk control, and vice versa.

4.2 *Risk*

4.2.1 The two components of the visualisation are the risk and reward/return that need to be linked together in a flexible, easily understood way.

4.2.2 We need to be able to answer (and display those answers visually) two questions:

- What is the level of risk of stocks, or groups of stocks?
- Which stocks, or groups of stocks, are most likely to move together?

4.2.3 Understanding how stocks, and groups of stocks, move together is vital to the process of portfolio construction, i.e. what is best way of risk managing a holding in Company A — is it to be underweight (or zero weighted or even 'short'), Company B in the same sector, or Company C in a different sector? In other words, we are trying to analyse which stocks cluster together in the market and which ones are very different to others?

4.2.4 The technique that we use to answer these questions is called cluster analysis — the visualisation of the results can be in the form of a dendrogram. Cluster analysis is a well-known statistical technique, and there is a rich literature using this branch of applied statistics in a wide range of the sciences (Everitt, 1974). The Appendix provides a simple introduction, but is not intended to be a comprehensive analytical guide.

4.2.5 Many of the clustering methodologies are not appropriate for investment, and are often not helpful for the purpose of risk management. Clearly, it is possible to cluster by a wide range of methodologies using a

wide range of metrics. However, they have to be useful and meaningful to the portfolio manager if it is to form a useful addition to portfolio risk management. Furthermore, it has to be consistent and integrated into the overall risk structure of the statistical factor model.

4.2.6 The method used is really very simple. Stocks are similar if they are close to each other in risk space — the closer they are to each other the more similar they are. The co-ordinates of an individual stock are the exposure of the stock to each factor. Using simple Pythagoras, we can determine the distance between two stocks, i and j , by the following formula:

$$\text{Distance} = \sqrt{\sum_{n=1}^N (f_i^n - f_j^n)^2}.$$

4.2.7 But what does this formula mean practically? Since the factors contained in the statistical factor model are orthogonal, the formula in ¶4.2.6 is the factor risk of an equally weighted long-short portfolio of stocks i and j ; clearly, this has some real investment meaning for an equity portfolio. This equation can be viewed as the factor distance. However, a further mathematical property of the statistical factor model is that the stock residual item is also orthogonal to each and every factor used in the model as well as being orthogonal to each other. Therefore, it is also possible to extend the factor distance by adding back the stock specific component of the risk model for stocks i and j , to derive the total distance measure. This is potentially an important extension, since the addition of the stock specific element may radically change the resultant shape of the dendrogram.

4.3 *Example of a Stock Dendrogram*

4.3.1 Typical results for the constituents of a pan-European index are shown in Figure 2. In this example we have used the fairly narrow FTSE €Stars Index. This index contains 29 stocks, a manageable number for this illustration. Other benchmark indices clearly have a different risk structure.

4.3.2 Figure 2 is built up from the left. The first thing to do is to find out which is the pair of stocks that are the most similar in terms of their risk, i.e. what two stocks have the lowest risk when one is hedged by the other. It is important to remember that all portfolios referenced against a benchmark index will be hedging long positions relative to the benchmark with short positions — they are just the same as a long/short hedge fund in this respect. Using just the underlying risk model, without knowing the identity of the stocks, the combinations of all pairs are calculated using the formula in ¶4.2.6, thereby identifying the pair with the lowest risk.

4.3.3 In this portfolio, the minimum risk ‘pair’ is ENI and Royal Dutch Petroleum — both of which are oil stocks; it is shown as Cluster D in




Figure 2. Dendrogram of the constituents of the FTSE €Stars Index

Figure 2. This means that the best way to hedge an overweight position (or long exposure for a hedge fund) in ENI with a single stock is to take up an underweight position (or short exposure for a hedge fund) in Royal Dutch Petroleum. The resulting (two-stock) portfolio would have a volatility of 3.7%.

4.3.4 As we move up the dendrogram, Cluster C (in this case a single stock) joins Cluster D. However, Cluster C is Total Fina Elf, so the combination of these two clusters forms an oil cluster.

4.3.5 Similarly, we can identify Dutch banking (Cluster B) and Spanish banking (Cluster A) groupings. This builds up further within the dendrogram to form a large bank-assurance cluster — a key part of the structure of this index.

4.3.6 On a broader basis, we can see that the FTSE €Stars Index naturally breaks down into two large clusters covering financials and telecommunications, media, technology (TMT) (Cluster E). Figure 2 shows that the TMT sector is, in systematic risk terms, very dissimilar from the financial sector, and since it only ‘joins’ with the financial sector at a very high level of total risk, it means that the two groups are not good hedges for each other. This index is fragmented into two large groupings. For asset

allocation purposes the most important — and risky decision — is the allocation between these two groups.

4.3.7 This is a simple example of how to visualise the risk dimension of stocks within a market and how they relate to each other in risk terms. Obviously, all of these comments are well known and well understood by portfolio managers. It is important to remember that, as the risk model is appropriate at a specific date (and is based solely on observable market returns), it will reflect changes in the risk structure of the market in an automatic and objective way.

4.3.8 However, for broader indices that are typically used as performance benchmarks the relationships between the individual stocks become more complex, and therefore more difficult to grasp intuitively:

- (1) The shape of the dendrogram changes over time as the structure of the market changes — there are clearly return generating opportunities as the market changes.
- (2) The generation of dendrograms can be repeated over time (a form of ‘dendrogram movie’ can be built up), and thus forms a consistent structured basis upon which portfolio managers can view the risk in their portfolios.

4.3.9 Figures 3 and 4 show the risk structure of the FTSE Eurotop Index for total risk and factor risk respectively.

4.3.10 There is a clear visual difference between the cluster results based upon total risk (Figure 3) and factor risk (Figure 4). The difference is attributable to stock specific risks and how they impact upon market structure. Obviously, all stocks join together at a higher level in the total risk structure than they do in the factor only risk structure. Less obviously, stocks with relatively high stock specific risks will cluster higher than their peers with lower stock specific risks. The total risk picture of market structure is more relevant for a concentrated portfolio (e.g. some types of hedge fund); on the other hand, if the portfolio is diversified (e.g. a style/ theme driven portfolio) the factor risk structure will be more relevant.

4.4 *Sector Risk Structure of the Market*

4.4.1 Conceptually, it is possible to do precisely the same analysis at the sector level to examine the sector risk structure of the market. Again, due to the mathematical properties of the risk model, it is possible to add up risk at any categorical level by calculating the weighted sum for each company in the category of each factor exposure. The stock specific elements can be added in if necessary. The formula to apply is:

$$\text{Factor exposure to category } i = \sum_{\text{all stocks } j, \text{ in category } i} w_j^i f_j^i.$$




Figure 3. Total risk structure of the FTSE Eurotop Index at the individual stock level




Figure 4. Factor risk structure of the FTSE Eurotop Index at the individual stock level

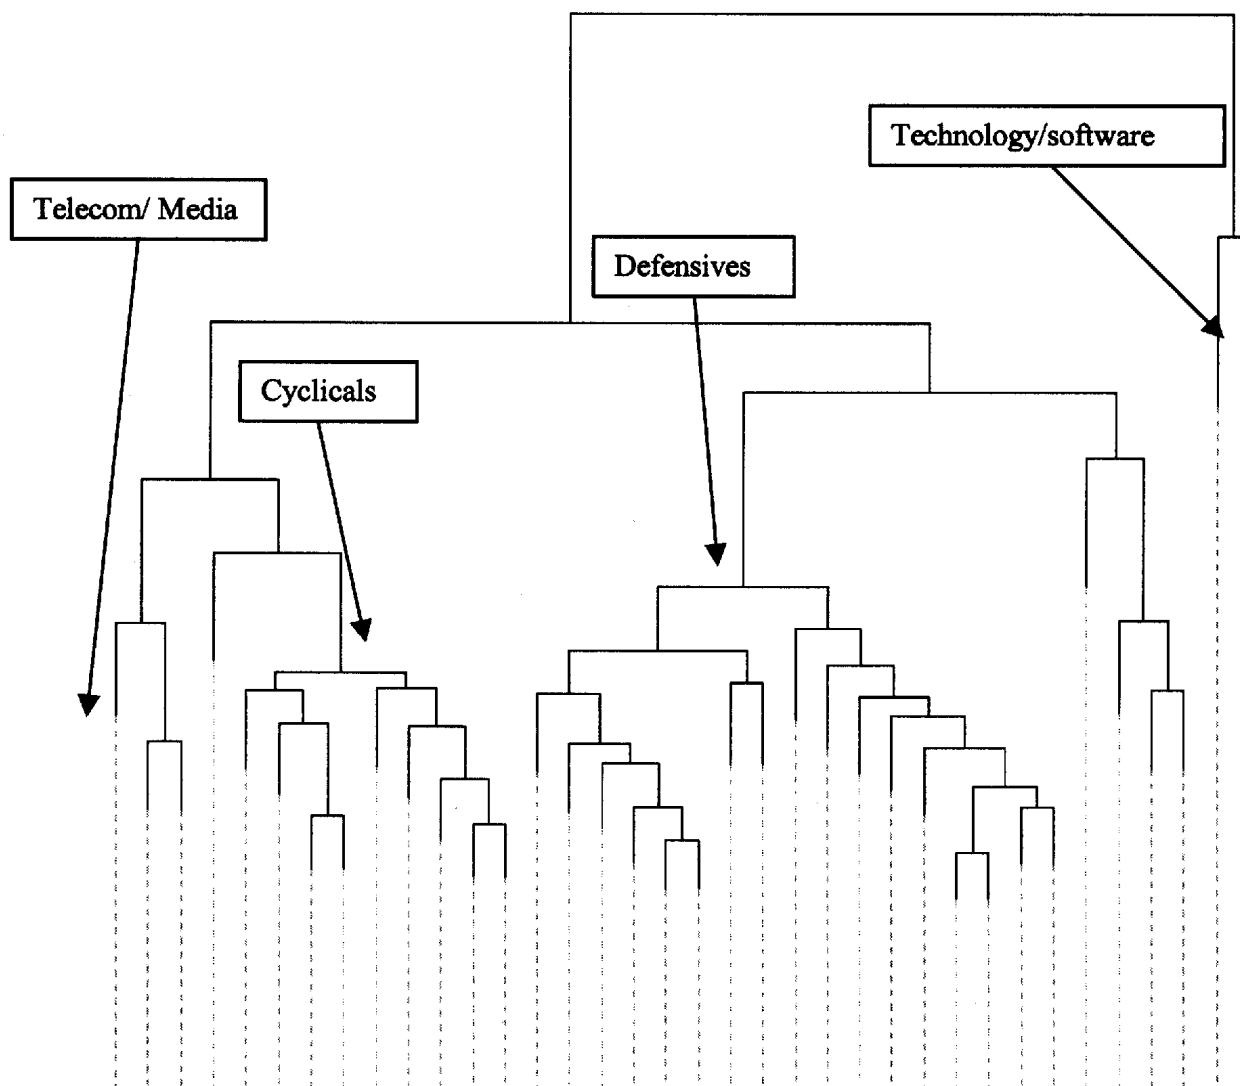


Figure 5. Factor risk structure of the FTSE Eurotop Index

4.4.2 As an example, Figure 5 shows the sector risk structure of the FTSE Eurotop Index at a particular point in time.

4.4.3 The dendrogram shows quite clearly the ‘high level’ sectors that were in force at the time — the technology/software group, the defensives and the cyclicals. It is interesting to drill down to look at the clustering at a lower level. Furthermore, the analysis can be carried on a systematic and repeatable basis to see how the risk structure of the market changes over time.

4.5 Reward/Return

4.5.1 Few portfolio managers have complete return forecasts on all of their holdings and the benchmark constituents. This makes the conventional approach to risk/return trade-off analysis very difficult. However, from our experience, most portfolio managers have a view on the returns that they are expecting from the stocks within their portfolios relative to each other.

4.5.2 Historically, this absence of comprehensive return forecasts has been a problem for risk management systems where an attempt is made to trade off risk and return. However, this problem can be partially overcome by using a risk statistic called marginal contribution to risk (MCR). The key features of the MCR are as follows:

- (1) The MCR can be positive or negative, and depends upon the stocks in the portfolio, in the benchmark and their relative weights.
- (2) If a portfolio manager wants to increase the risk in a portfolio, then he can either add to stocks with a positive MCR and/or sell stocks with a negative MCR.
- (3) It can be driven by beta, in the sense that in a portfolio with beta greater than one all stocks with positive beta will tend to have a positive MCR, because adding to these stocks will increase portfolio beta and hence risk.
- (4) Stocks with a negative (or relatively low magnitude) MCR are diversifying the risk within the portfolio.
- (5) In an unconstrained portfolio that properly reflects and captures the views of the portfolio manager, the MCR is proportional to the expected return on the stock, i.e. for two stocks with different MCRs, the stock with the higher MCR will have a higher expected rate of return forecast.

4.5.3 Therefore, the MCR forms the bridge between risk and reward.

4.6 *A Worked Example*

4.6.1 In order to see how this works in practice, we have constructed a simple portfolio analysis using the FTSE Eurotop 300 as the benchmark and the FTSE €Stars Index as the portfolio.

4.6.2 If a portfolio were constructed based upon the €Stars Index to track the Eurotop 300, it would have an expected tracking error of over 7½%. By any institutional definition this would be deemed an aggressively run portfolio. However, incurring risk is not necessarily a bad feature in a portfolio, provided that the portfolio has been constructed efficiently and the risk budget has been spent appropriately. Table 2 shows the main risk characteristics of the portfolio.

4.6.3 Table 3 shows the ‘top’ and ‘bottom’ ten marginal contributions to active factor risk from the stocks within the portfolio — the stocks in the middle have been omitted for convenience.

Table 2. Summary risk statistics of FTSE €Stars vs. FTSE Eurotop 300

Tracking error	7.61%
Portfolio beta	1.14
Correlation	0.93

Table 3. Holdings listed by descending marginal contribution to active risk

Company name	Marginal contribution to active risk	Company name	Marginal contribution to active risk
Top 10		Bottom 10	
Nokia	35.55%	Bayer AG	8.60%
Alcatel	28.67%	Generali Assicurazioni	7.91%
Koninklijke Philips Electronic	23.53%	Daimlerchrysler	7.56%
Siemens AG	23.24%	Royal Dutch Petroleum	7.14%
Vivendi Universal	22.62%	Ente Nazionale Indrocarburi (ENI)	5.94%
Deutsche Telekom	22.21%	Carrefour	5.21%
Telecom Italia	18.01%	L'Oreal	4.34%
Telefonica	17.28%	Suez	4.19%
AXA	16.47%	E.ON	1.23%
Allianz AG	15.71%	Unilever NV CVA	−5.15%

4.6.4 In rate of return terms, it is implicit in the portfolio that the manager is expecting the return on Nokia to be greater than on Alcatel, which, in turn, will be greater than on Philips, and so on. If this is not the case, then the portfolio has not been constructed efficiently.

4.6.5 Figure 6 shows the dendrogram for the benchmark. The shaded bars indicate the disposition of the portfolio across the benchmark — different colours represent the sector disposition of the portfolio.

4.6.6 The marginal contribution to risk is shown at the foot of the page. Those stocks where the MCR is negative are shown in grey — the expected tracking error can be reduced by buying/increasing the weight in these stocks. The coloured bars represent the MCR for the stocks that are actually held, based on their actual weights.

4.7 *Observations on the Structure of the Dendrogram and the MCR*

4.7.1 The dendrogram represents the risk structure of the benchmark.

4.7.2 As the market changes, so the structure of the dendrogram changes. These changes can take place at the ‘macro’ or at the ‘micro’ level. Examples include:

- (1) style rotation, e.g. value versus growth;
- (2) investment themes (the evolution of the TMT phenomenon and then its partial demise); and
- (3) Spanish banks form their own cluster, independent of the other continental European banks, presumably due to their heavy exposure to Latin America.

There is very clear evidence of sensible market-related clustering.




Figure 6. Total risk dendrogram, overlaid with the portfolio and the MCR

4.7.3 From Figure 6, there is a visual impression of stock, sector and (in some instances) country clustering, how the clusters join together and, importantly, how high the clusters are when they join together. For example, the media and technology cluster is quite separate from the rest of the market (in investment terms, it means that it is difficult to hedge with stocks from other clusters), and it is more ‘risky’ than the other main parts of the market.

4.7.4 The majority of large telecom stocks is quite separate from the main clusters of media and technology stocks.

4.7.5 Some individual stocks are quite different to anything else in the market.

4.7.6 The MCRs are typically positive, primarily because the beta of the portfolio is greater than 1 — see Table 2. This implies a positive expected return on the market as a whole.

4.8 *Observations on Portfolio Construction and Risk Budgeting*

4.8.1 The portfolio is not well diversified. The coloured bars indicating the portfolio holdings in Figure 6 are not spread out across the clusters — they are relatively concentrated. In order to reduce the tracking error, stocks in other clusters would need to be selected. This can be analysed in more depth and more properly in the section on risk budgeting.

4.8.2 The portfolio does not contain any U.K. stocks, despite heavy exposure in the benchmark — this is a cause for lack of diversification.

4.8.3 There are several important clusters where there is no portfolio exposure. These correspond to stocks with large negative contributions to MCRs.

5. CHANGING RISK STRUCTURE OF AN EQUITY MARKET

5.1 *Risk Changes over Time*

5.1.1 Market practitioners all know that the risk structure of equity markets change over time. In order to frame an investment view, it is imperative that portfolio managers know, on a disciplined basis, both the current risk structure of the market and how it has changed. Only when this knowledge is available is there any chance of taking rational investment decisions.

5.1.2 The statistical approach to building risk models is particularly well suited to an analytical approach to assessing the current risk structure of an equity market. As we have discussed in Section 3, risk models can be built over various time periods, with various frequency of observations (e.g. daily or weekly rate of return measurements), so that risk can be assessed over a whole variety of periods that range from the short-term to the long-term. Accordingly, the resultant risk models can be used for trading, hedge funds or longer-term investors.

5.1.3 In Section 4 we saw how we could use a risk model to see the risk structure of the market on a consistent basis. In this section we develop a methodology to examine how, if at all, the risk structure of an equity market changes over time.

5.1.4 We are interested in determining if several dendrograms generating different data sets describe the same classification. The main idea underlying our approach is to analyse the similarity of dendrograms by studying how similar the underlying distance matrices are. In our case, we have different dendrograms for the same stocks generated by using returns data over different periods.

5.1.5 In its simplest case, we have used Spearman's rank correlation coefficient ρ as the measure to test for similarity. Despite its slightly unfamiliar looking formula, ρ is just the correlation calculated for ranked vectors — in our case the ranking of the distance matrix, considered as a vector. Therefore ρ takes values between -1 and $+1$. It takes the value $+1$ when the two ranked vectors are identical, -1 when the rankings are in opposite order, and small absolute values are taken when the two vectors are unrelated. Under the null hypothesis that the two vectors have been randomly drawn from some population, a function of ρ is distributed as Student's t .

5.1.6 Let X and Y be two vectors of dimension N . Let d_i be the difference between the ranks of the i th entry of X and Y . Then ρ is given by:

$$\rho = 1 - \frac{6 \sum_{i=1}^N d_i^2}{N(N^2 - 1)}.$$

5.1.7 In Section 4.4 we observed that it was possible to aggregate risk across categories so that, for example, we could analyse the risk structure of the market by sector. In a similar way, we can aggregate risk in the same 'category', and apply the results to the Spearman coefficient in ¶5.1.6. This enables us to investigate how the risk structure of the market is changing at, for example, the industry sector level of a market.

5.2 A Worked Example

5.2.1 In our example, we will look at how the risk structure of the European market has changed during the five-year period from January 1998 to December 2002. This is a particularly interesting period, since it spans the formation of the single currency in Europe as well as the rise and fall of the 'internet boom'.

5.2.2 In terms of the introduction of the euro, we would expect to see the reduction of the country effect, i.e. the risk structure of the market would become more stable when looked at from the perspective of the country aggregation.

5.2.3 Similarly, if the methodology is to be of any use at all, then, at the very least, we should be able to quantify and see the effects of the ‘technology/media/telecom’ phenomenon, both in terms of the emergence of the grouping and the subsequent changes that took place in the first quarter of 2000.

5.2.4 We have looked at a broad universe of over 600 European stocks, including the U.K. The first requirement is to build risk models at each month end of the stocks that were in the market at that time. We used the historic constituents of the DJ Stoxx Europe 600 Index. One advantage of this is that the historic database allows for the survivorship database in that it shows the actual index constituents at each historic point of time. However, some corrections were made where necessary, to allow for changes in the definition of industry sectors and groupings where appropriate.

5.3 *The Country Effect*

5.3.1 Based on the risk models built, it is a relatively easy task to calculate the Spearman coefficients for any pair of start and end dates within the five-year observation period. If both the start and end dates are the same, then the Spearman coefficient is +1. To the extent that the statistic for any pair of dates is different to one, it shows how the risk structure has changed between the two dates. Obviously our models are estimated on overlapping data sets through a moving window, and this has some implications for the statistical power of these tests and their interpretation.

5.3.2 The Spearman rank correlation coefficients are determined for all pairs of start and end dates over the observation period. The results are then put into a symmetric matrix, with one axis the start date and the other axis the end date. Rather than display these results as a table, it is preferable to display the results visually in the form of a heat map.

5.3.3 Figure 7 displays the results. The colour code goes from yellow (representing a Spearman coefficient of 0) up to red (representing a Spearman coefficient of +1). The more red there is in the heat map the more similar the dendrograms are between the start and end dates. By definition, the leading diagonal will be entirely red. The right hand panel of Figure 7 shows the Spearman coefficients in adjacent periods.

5.3.4 It is very clear from Figure 7 that the risk structure in Europe from the perspective of country has been remarkably stable since the middle of 1999. Prior to that there was more variability. (If the analysis is undertaken over a longer period, there is also very clear evidence of variability in the country risk structure.)

5.4 *The Sector Effect*

5.4.1 If we aggregate the risk by industry sector and repeat the analysis in Section 5.3, we get a visualisation of how the risk structure of the market, by industry sector, has been changing. The results are shown in Figure 8.

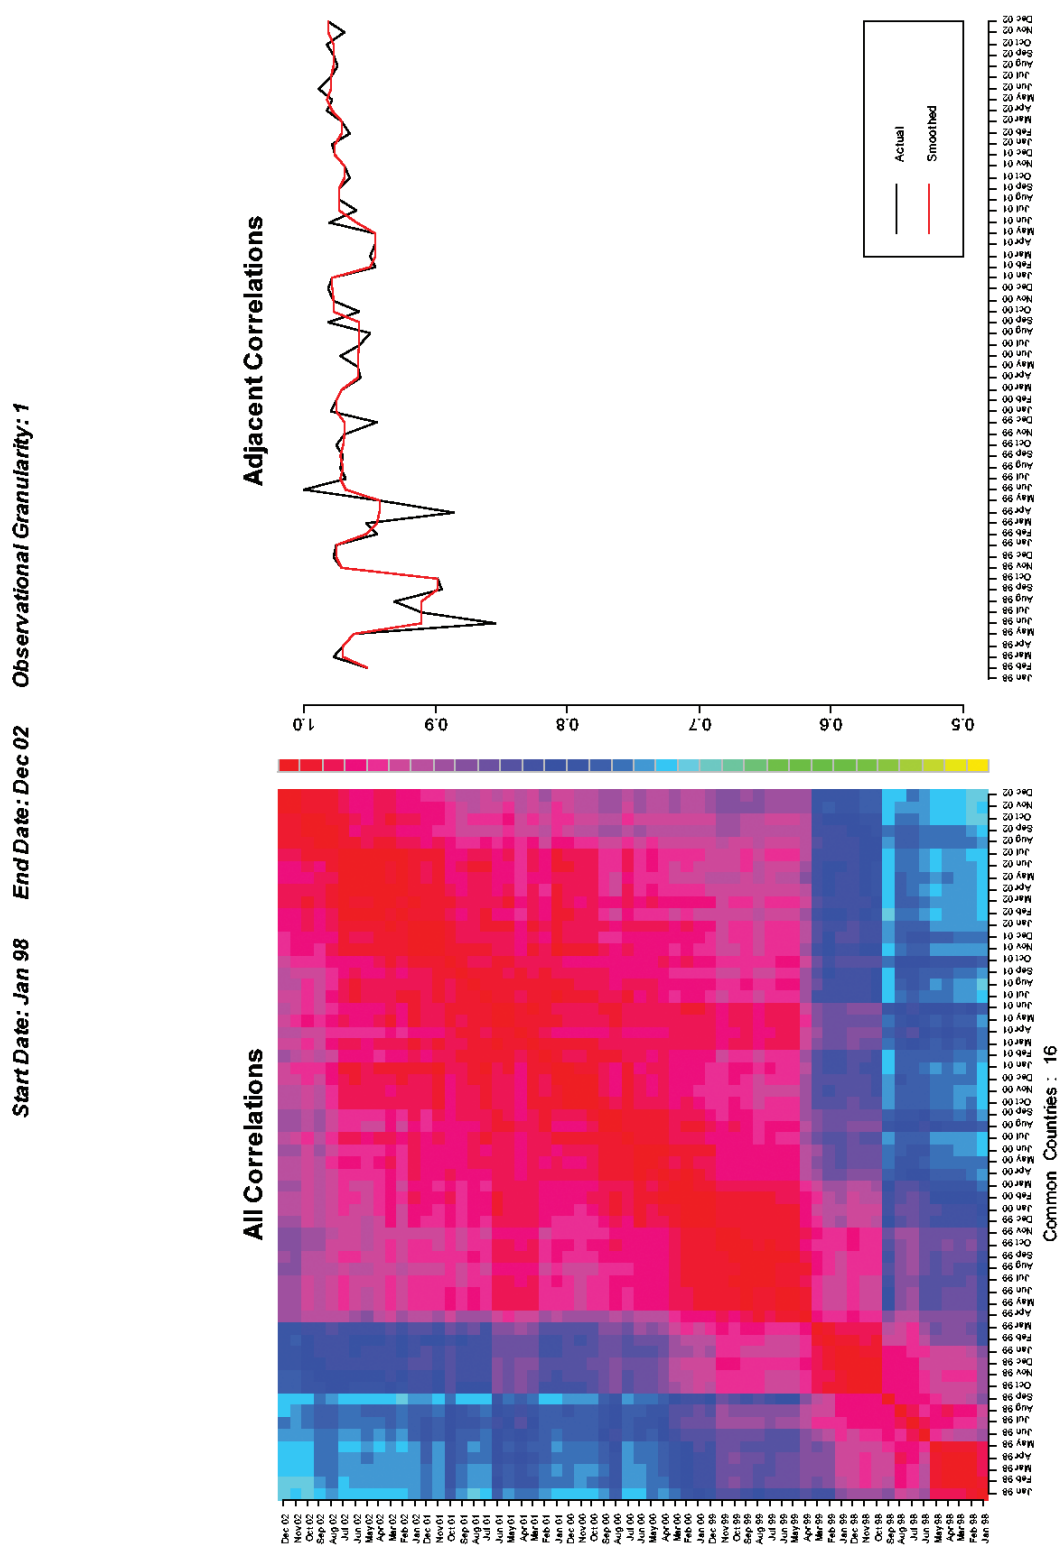


Figure 7. Country risk structure of Europe 1998-2002

Start Date: Jan 98 End Date: Dec 02 Observational Granularity: 1

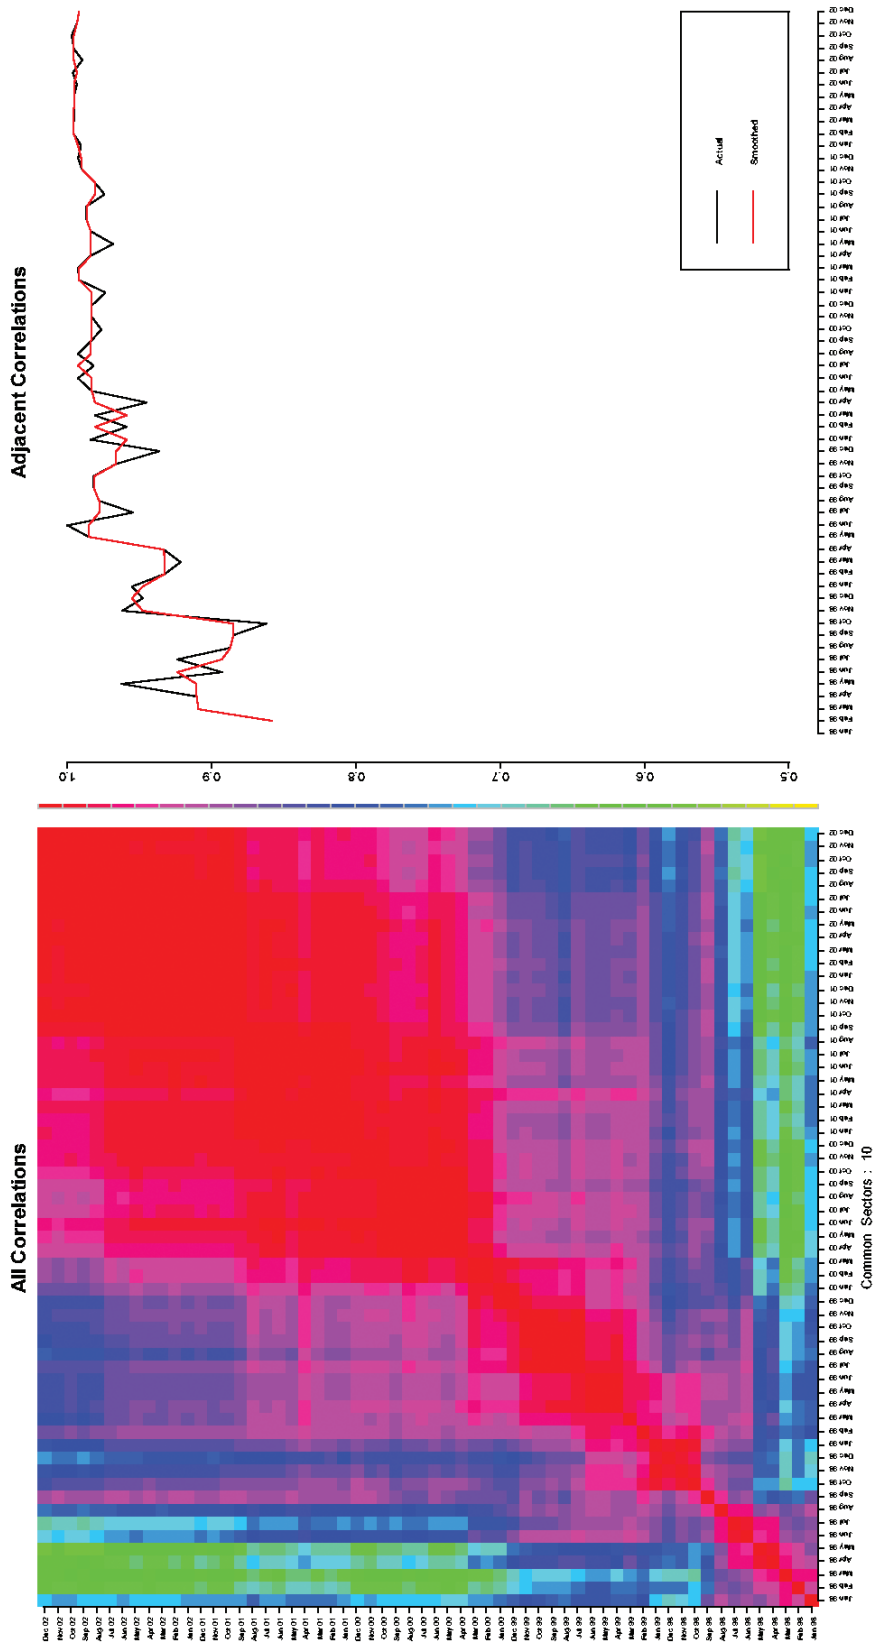


Figure 8. Industry risk structure of Europe 1998-2002

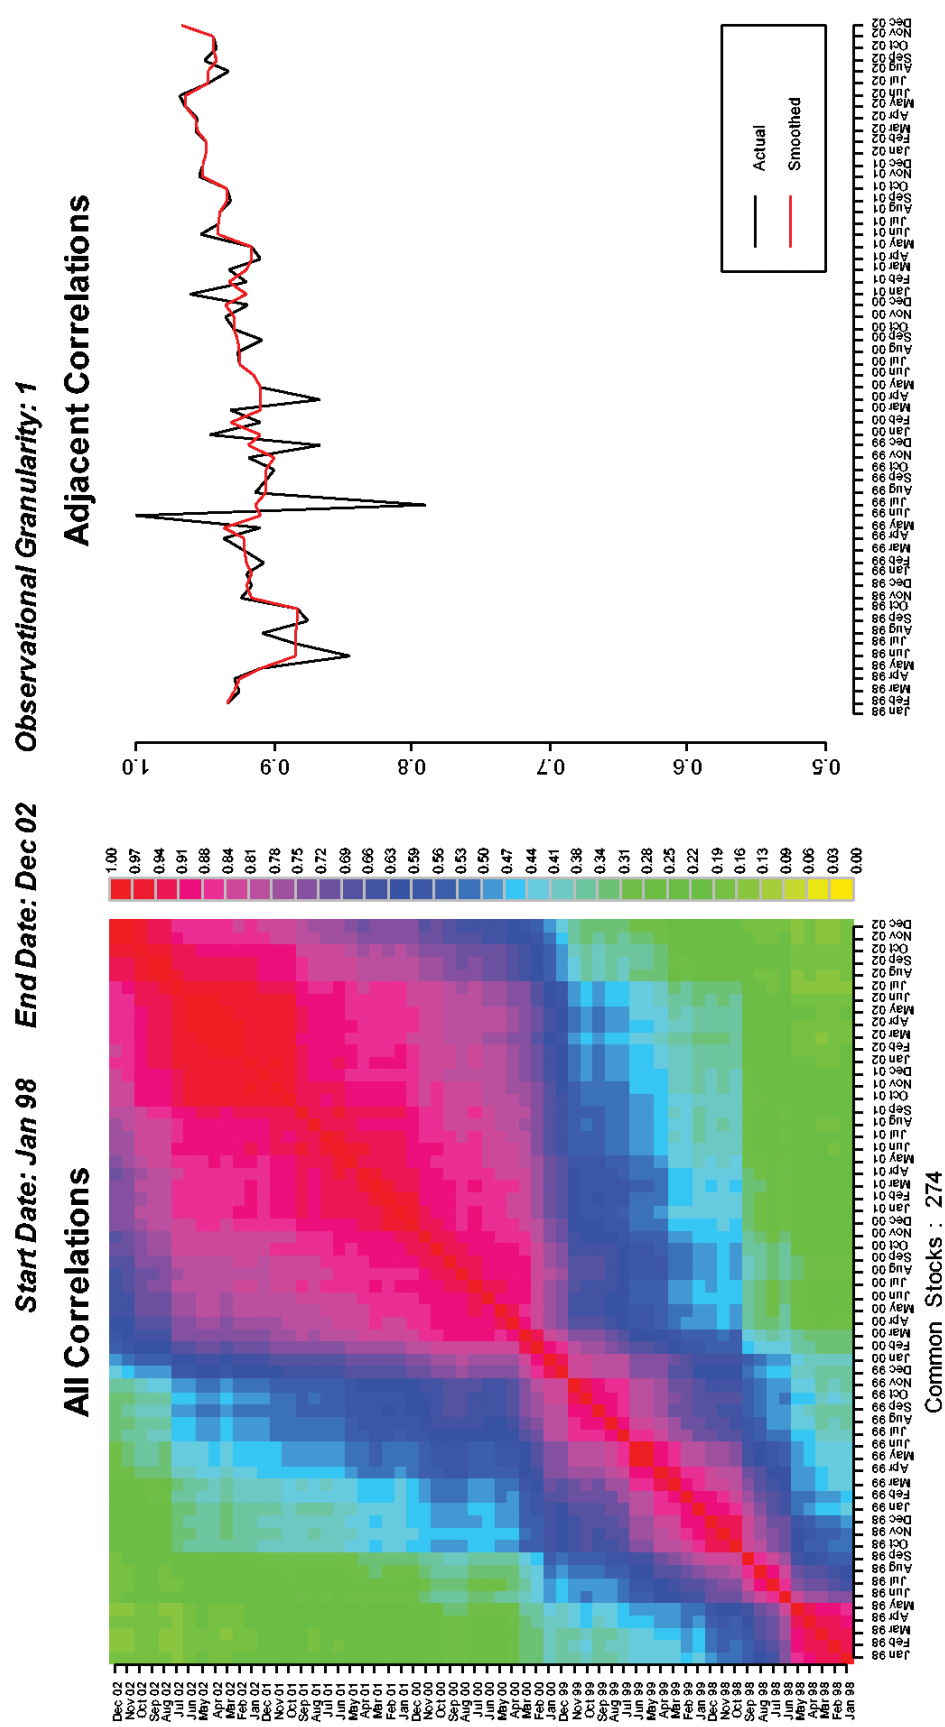


Figure 9. Stock risk structure of Europe 1998-2002

5.4.2 Figure 8 shows significant differences in the industry risk structure to that of countries. There is an abrupt change of colour almost mid-way through the period — in fact March/April 2000. Prior to then there was a lot of change in sector risk structure. Since then, and particularly during the period August 2001 to December 2002, the structure has been relatively stable.

5.4.3 The early part of the analysis period reveals a lot of sector rotation — this is particularly evident by looking at the right hand panel of Figure 8. This shows that the Spearman coefficients between adjacent periods have fluctuated quite wildly.

5.4.4 There are large differences between the industry risk structure at the start and end of the period.

5.5 *The Stock Effect*

5.5.1 Finally, a similar analysis can be completed for stocks. The results are shown in Figure 9.

5.5.2 It is little surprise that the results show far more changes in the risk structure of stocks from one period to another and that the overall level of the Spearman coefficients are lower than for both countries and industry sectors.

5.5.3 There are modifications that can be made to the approach outlined in this section. Since we have shown how easy it is to aggregate risk due to the mathematical properties of the statistical factor model, it is possible to calculate distance measures between countries, industry groups, sectors, etc. (or any other categorisation) at a single point in time, and then build these up over time to visualise how the market changes over time.

5.5.4 Some specimen results are shown in Figure 10. For illustrative purposes, we have aggregated the risk model by industry group and show the results as at the end of each year 1997 to 2002. More frequent observations would show the dynamic effects of the change in the market through time. Figure 10 represents the risk structure with an appropriate lag. Instantaneous risk views could be obtained by using shorter-term models, as explained in ¶3.5.7.

5.5.5 Again, we present the results in the form of a series of heat maps. No clear picture emerges at the industry group level from the pictures shown as at December 1997 and December 1998. The technology and telecom industry groups are very clearly in evidence from the pictures at the other observation points. The energy sector shows a similar characteristic.

5.5.6 It is clearly possible to aggregate at the industry or sub-industry level to see a more detailed breakdown of the structure of the equity market.

5.5.7 The visual results that we have presented clearly demonstrate the existence of structural changes associated with the euro convergence




Figure 10. Industry risk structure of Europe, 1998-2002






Figure 10 (continued). Industry risk structure of Europe, 1998-2002

Weighted Averaged Sectors



Weighted Averaged Sectors




Figure 10 (continued). Industry risk structure of Europe, 1998-2002

phenomenon and the TMT bubble. It is worth pointing out that one can explicitly test for structural breaks using a variety of parametric and non-parametric statistical tests to give a level of significance.

5.5.8 From a practitioners point of view, this is potentially important. Changes in market structure could be viewed as either temporary (and mean reverting) or reflecting some structural change.

6. RISK BUDGETING

6.1 *The Need For Risk Budgeting*

6.1.1 Risk budgeting has received a great deal of interest from the investment management community recently (Rahl, 2002), but no clear consensus has emerged on how it should be implemented. Many of the ideas in this section could equally be referred to as the decomposition of expected risks or ‘co-variance accounting’. Risk budgeting takes the analysis further: “Where do we want to take risk; where do we want to avoid it, whilst satisfying various constraints?” These types of questions can only be answered in conjunction with an alpha generation model (whether quantitative or otherwise) and with some form of simulation. The vital point is that a disciplined and repeatable framework for accounting for risk is a necessary pre-requisite.

6.1.2 In a traditional, long-only investment environment, there is a natural limit to the downside risk, which is the case where the value of the asset falls to zero. In a long-only environment, no position can turn into a liability except in relation to the benchmark. A benchmarked long-only portfolio is merely comprised of:

- (1) the benchmark index; plus
- (2) a long-short overlay portfolio with two constraints: (a) a net zero exposure by market value; and (b) no short position may be in excess of the benchmark index exposure.

These constraints are an arbitrary restriction on the long-short overlay portfolio. When they are lifted, as in the more general case of a hedge fund, there is a pressing need for some new discipline to be imposed on positions and risk.

6.1.3 Another environment that raises the need for risk management is where a portfolio is split into sub-portfolios that are managed by different people. The actions of the individual managers need to be coordinated, so that the portfolio as a whole is not exposed to inappropriate levels of risk (Sharpe, 2002).

6.1.4 In both these situations, there is a pressing need to manage and control the sources of risk. The traditional business method for making such allocations is to budget and control. In the context of risk, efforts to budget

in an analogous way have been only partially successful. In this section we will explore some of the issues involved in practical risk budgeting.

6.1.5 A budget is a meaningful and useful method of allocating resources available to a manager, if the manager then has some freedom and discretion in choosing autonomously how the budget is to be spent. Part of this process will be to find the ‘goods’ (in our case returns) that have the cheapest price (in our case, contribution to portfolio risk). The complication in the case of risk, and in particular the case of systematic risk, is that the decisions of one manager affect the prices faced by another.

6.1.6 The way in which this works is clear from the mathematical expression for portfolio risk, which consists of terms like $w_i w_j \sigma_{ij}$. Most terms have two decision variables (w_i and w_j) implicated, so attributing responsibility for the size of the term $w_i w_j \sigma_{ij}$ is difficult. If manager i has control over w_i , and manager j has control over w_j , the size of the typical term $w_i w_j \sigma_{ij}$ is determined jointly between the two managers. The exception is terms where $i = j$, which are terms in variance and not covariance — they are specific risks, and can be budgeted for in a fairly conventional way.

6.1.7 The remaining terms with $i \neq j$ are contributed by systematic risk. For the large portfolios of many assets that are typical of institutional investments, it is normal for the great majority of the total risks to be systematic. Active risk relative to the benchmark typically has a lower proportion of systematic risk, but it is not at all unusual for it still to dominate specific risk. Hence, the management of systematic risk is a matter of considerable importance, but, as we have seen, it is not obviously amenable to budgeting in a conventional way.

6.1.8 We will argue that, at the moment, we are better equipped to account for risk than we are to budget it. Because, in a portfolio risk context, the terms $w_i w_j \sigma_{ij}$ cannot be uniquely attributed to one manager, they are jointly owned by manager i and manager j . There is a need for an overview of all managers’ positions by a person with responsibility for the entire portfolio. Budgeting becomes an interactive process, in which the responsibility for the covariance terms is determined in a way that is adapted to suit the investment process.

6.2 *Covariance Accounting From The Bottom Up*

6.2.1 There are a number of alternative models for risk, and earlier, in Section 3, we discussed the relative merits of these alternatives. The same principles of accounting can be applied to any multi-factor model for expected risk. The first assumption that we make is that forward looking risk is quantified in terms of variance of returns. In general, these returns can be either absolute returns or returns relative to a benchmark, but, for the sake of brevity, we will henceforth assume that we are looking at relative returns, and use the term ‘tracking error’ to describe the square root of risk variance.

6.2.2 Our second assumption is that risk is best understood at the level of the individual security. A collection of securities, whether it be an index or a portfolio, will be better analysed and predicted, both in terms of risk and return, if the changing weightings of individual securities are taken into account.

6.2.3 The objective that we have set ourselves is to account for risk; literally to attribute each basis point of expected tracking error to its sources. The calculations of risk take into account the positions in each constituent security — we do not know of another way to analyse risk in a consistent way which will yield the same result at different levels of aggregation.

6.2.4 As we have already noted, the expression for variance is quadratic; a summation of covariance terms of the form $w_i w_j \sigma_{ij}$:

$$\sigma_p^2 = \sum_{i,j} w_i w_j \sigma_{ij}.$$

6.2.5 The most natural way to arrange these is in a square table, with the dimension on each side being the total number of securities. The summation of the terms in this table can be done either all at once or in a hierarchy of levels (e.g. by industry, then by sector). Regardless of the level of aggregation at which we view the analysis, the total figure is the same.

6.2.6 The total figure that remains the same is the figure for variance. The algebra of variance is relatively simple, and we can add it up in the order that we find most useful. The trouble with variance is that the units are, for most people, unfamiliar, and therefore difficult to interpret. The conventional solution is that the square root is taken. We take a different route, and divide each term by σ_p , the portfolio volatility. The result is that our attribution of covariance is now in units of volatility.

6.2.7 We now have a tabulation of the sources of risk, which can be aggregated at any level or in any order, and which always adds to the tracking error or total risk. Each element is of the form $w_i w_j \sigma_{ij} / \sigma_p$:

$$\sigma_p = \frac{1}{\sigma_p} \sum_{i,j} w_i w_j \sigma_{ij}.$$

6.2.8 This means that the term is zero if the active weight in either asset is zero, the contribution is zero, and if the covariance is zero, the contribution is also zero. Thus the attribution works, both from an algebraic point of view and from a commonsense point of view.

6.3 *Risk Factors*

6.3.1 If the bottom line for risk can be attributed between the positions in different assets, it can also be attributed between categories of risk. As we

noted before, there are many competing multi-factor models for risk, most of which conform to a simple three-class categorisation, which we adopt here.

6.3.2 The first category is specific risk; the part of a security's returns which is totally unrelated to other securities' returns. The second category is 'market risk', in the sense meant by the capital asset pricing model — that is, the sensitivity of a security's returns to the returns on a broad-based market index, as measured by beta. The third category is the remaining sources of systematic risk, which we will call 'styles and themes', which include what is normally known as industry risk, as well as size, value, and so on.

6.3.3 We can relate this to the risk account as follows. As we saw before, each term in the account looks like $w_i w_j \sigma_{ij} / \sigma_p$. In a multi-factor risk model, the covariance σ_{ij} is broken down into its sources. For our purposes this does not present a problem, because we can write each term as the sum of three components:

$$\sigma_p = \frac{1}{\sigma_p} \sum_{k=1}^3 \sum_{i,j} w_i w_j \sigma_{ijk}.$$

When: $k = 1$, σ_{ijk} corresponds to the stock-specific risk;
 $k = 2$, σ_{ijk} corresponds to the market risk; and
 $k = 3$, σ_{ijk} corresponds to the style and theme risk.

6.3.4 The third category can, of course, be broken down further. The point here is simply that the distributive law applies, so if we have an additive multi-factor model for covariance, then we also have an additive multi-factor method of accounting for risk. This simple additive property makes for easy accounting, and, as we shall see, easy visualisation.

6.4 *Aggregations of Risk*

6.4.1 As we noted earlier, there is a fundamental assumption that the best method for measuring and analysing risk is at the security level. In this way we take account of the current composition of a portfolio, a benchmark, or any other aggregate such as an industry or country, and not the historical composition that varied over time and is now out of date. Analysing risk at the security level is not a problem, given the computing power that we have available. The challenge is to display and communicate the information in a way that is intelligible to the portfolio manager.

6.4.2 The problem of information overload is well documented in the field of accounting. When presented with just a simplified set of accounts, the average analysts are more likely to pick out the key points than if they are presented with a detailed set of accounts. There is a pressing need to simplify and clarify, and to present detail only where and when the user requires it.

One solution is to present the information in terms of aggregates, and to allow the user the possibility to look into these aggregates on demand. This functionality is commonly referred to as ‘drill down’.

6.4.3 The process of aggregation is simply a way of adding up the type of terms that we saw before — $w_i w_j \sigma_{ij} / \sigma_p$. The tracking error is a sum of all of these elements over the entire range of i and j , which, for an investment universe of 700 securities, is a square table composed of 490,000 entries. It is not helpful to present the user with the stark choice between a single number or a table with 490,000 entries. What is needed is a system of classification of the securities, within which we can aggregate the risk meaningfully. Each security must belong to exactly one category. One such system of classification is the sector classification, another would be country classification, and a third would be classification using statistical cluster analysis.

6.4.4 The aggregation process consists of adding all the terms that belong to a particular category, for example a sector. As an example, say we wish to calculate the portion of tracking error that is attributable to the positions in the oil sector interacting with the telecoms sector. We do the summation of terms $w_i w_j \sigma_{ij} / \sigma_p$ for securities i which are members of the oil sector and for securities j which are members of the telecoms sector. Equally, we can use different categorisation schemes on the two axes, so, for example, in a pan-Euro portfolio we could calculate the tracking error which is attributable to positions in Germany interacting with positions in the pan-Euro technology sector.

6.4.5 The aggregation process leads to a very similar matrix to the one that we had originally, but one with a smaller number of elements. The units of the resulting matrix are still tracking error, and they still add up to the total tracking error. What we have described is the bottom-up accumulation of tracking error into useable high-level aggregates.

6.5 *Displays and their Interpretation*

6.5.1 The basic element of risk in all preceding sections has been a term like $w_i w_j \sigma_{ij} / \sigma_p$. In contrast to ordinary accounting, where the numbers are most naturally presented in a column, the elements of risk are most naturally arranged in a table. One could call this table a matrix, but since most matrix operations are not appropriate, and the only natural operations to apply to it are to add up the elements, it is more helpful to think of it as a table.

6.5.2 We are most interested in elements in the table that are large in magnitude, which either add a lot of risk to the portfolio or remove a lot of risk from it. The size of the table can be large, so what is needed is a method of display that assists with the process of identifying the major contributions to risk. One method is to order the elements so that the most important are all grouped together. We group the elements of greatest importance at the top left, with the elements of lowest importance at the

bottom right. Another method uses colour codes to produce a ‘heat map’, rather like a night vision camera uses false colour to highlight animate objects. The colours are selected to highlight the largest values, so that there is an instant perception of the key sources of risk — the ‘hot spots’.

6.5.3 As we noted before, the table can be analysed into a set of three (or more) tables which add up to the total table, and which correspond to the different risk factors. The specific risk table shows the part attributable to unsystematic, stock-specific risk. The ‘market risk’ table shows how the betas of the assets add to or reduce the portfolio risk, and will help to highlight the positions that have the largest impact on the portfolio’s active beta. The ‘styles and themes’ table shows which assets contribute most to other sources of systematic risk, purged of beta. These tables and their alternative visual representations as ‘heat maps’ are the basic tools of risk accounting.

6.6 *Risk, Marginal Contribution, and Expected Return*

6.6.1 Thus far we have looked at accounting for the expected risk in a portfolio, whether it be tracking error or total risk. Another analysis which can be useful is the marginal risk, that is the sensitivity of risk to the weighting in security i .

6.6.2 The expression for marginal contribution to risk (MCR) for security i is a single summation of terms $w_j \sigma_{ij}$:

$$MCR_i = \frac{1}{\sigma_p} \sum_j w_j \sigma_{ij}.$$

6.6.3 The constituent terms in this summation tell us how the positions in other, correlated assets contribute to making this security risky, at the margin. Once again the units can be made more helpful by dividing through by the portfolio risk, so that the marginal contributions are marginal contributions to volatility rather than to variance.

6.6.4 In the area of marginal contribution and risk budgeting, others have at times made the assumption that the portfolio is optimal. We consider this to be an assumption that is, at best, heroic. Next, to make useful progress in the analysis, others have assumed that there are no binding constraints on the portfolio. This is an assumption that is simply untrue in the majority of practical situations.

6.6.5 Without these assumptions, there is still much that can be inferred from an inspection of the MCR. MCRs of assets which are not bound by a constraint can be directly compared; if the rankings of their MCRs are not in accordance with the managers’ rankings of their expected returns, then there may be an opportunity to add value by trading, depending on the costs of trading.

6.7 *Risk Budgeting and Risk Management*

6.7.1 Budgeting is an exercise in containment, normally applied to the containment of costs. The motivation to devise an analogous process to contain risk in a portfolio is strong. However, applying budgetary control to portfolio risk is problematical, because the systematic risk can build up in a portfolio through the interactions of its constituents, and these interactions, by their nature, are not readily contained. Specific risk can be budgeted for, because these interactions do not occur, but only in special cases is this the dominant source of risk.

6.7.2 One way of getting around this is to make assumptions, even assumptions that are implausible, such as zero correlation between different managers' active returns. When the data are available to analyse risk without resorting to assumptions, the exercise of due diligence requires use of these data. We therefore reject the use of assumptions, be they plausible or not.

6.7.3 The original problem that motivated risk budgeting was the containment of risk. We propose a risk management process that satisfies this need for containment of tracking error at the portfolio level, without requiring the implausible assumptions. An effective process places some demands upon both the information systems and the operations of a portfolio management business.

6.8 *A Worked Example*

6.8.1 For the sake of concreteness, consider the example of a global portfolio, with separate managers for different regions. The following are some suggestions for an effective risk management process:

- (1) *Management of variance.* The individual portfolio managers have a budget for the risk taken within their portfolios, without reference to other portfolios, i.e. the variance. Staying within this budget is the responsibility of the managers, and they can ensure compliance by reviewing the impact of trades before executing them. Market moves in their own sub-portfolio and in the rest of the portfolio will have an impact on portfolio weights, and hence on their risk contribution, so a periodic review is also needed.
- (2) *Management of covariance.* As we have pointed out, a significant part of the total risk is typically contributed by covariance. The active positions of the Pan-Euro manager are very likely correlated with those of the North America manager, but the responsibility for this covariance cannot be uniquely assigned to either manager. A solution is to assign half the covariance to each manager, and to give each manager a budget for covariance. Staying within this budget is the responsibility of the manager. If a trade by the Pan-Euro manager would breach that manager's covariance budget, or anyone else's, it would need approval by a person with responsibility for the portfolio as a whole.

- (3) *Interactive display of variance/covariance accounts.* The management of risk is a multi-level affair. What applies to management of risk accumulating between regional sub-portfolios also applies to the management of risk accumulating between the sectors or countries within a portfolio. Because the risk contribution tables are potentially large, even at the industry/sector level, the use of well-designed information display technology makes a big difference. Management by exception highlights clearly the biggest contributors, and, in particular, those that are outside their control range. Drill down allows the user to move from an aggregated level to a detail level and back again.

6.8.2 In this example we will analyse the risk budgeting process of a global portfolio apportioned by region — North America, Europe (including the U.K.), Japan and the Pacific Rim. In each region the largest market capitalisation stocks (in descending order of market capitalisation) are held, such that they represent the top 30% of each region's market capitalisation. In this way we have constructed a portfolio that has a large capitalisation bias.

6.8.3 *Summary risk statistics*

The first analysis is to calculate the summary risk statistics of the portfolio. These are summarised in Table 4. The expected tracking error is 4.91%. The purpose of the risk budgeting is to analyse and slice and dice this tracking error of 491bps in a variety of ways, and to drill down so that the sources of the risk can be determined.

6.8.4 *Risk budgeting by country*

- (1) Table 5 shows the apportionment of the 491bps by country. It reveals both the risk incurred within each country and, most importantly, the effect of the co-variance terms between countries.
- (2) For ease of reading we have excluded the entries in the table that contribute less than ± 1 bp to the total.
- (3) The table of results can also be represented in a heat map, so that the relative importance of the contributions to the tracking error can be easily seen. The results are shown in Figure 11.
- (4) It is clear from Figure 11 that most of the risk is being taken in the United States of America, with lesser amounts in the U.K., Japan and Finland. There is little diversification across country — the largest contribution coming from the co-variance term between Germany and the U.K.

Table 4. Summary risk statistics

Tracking error	4.91%
Portfolio beta	1.00
Correlation	0.96

Table 5. Risk budgeting — by country

	United States	United Kingdom	France	Netherlands	Japan	Switzerland	Germany	Spain	Sweden	Italy	Australia	Belgium	Singapore	Greece	Canada	Finland	TOTAL
United States	305	29		5	(2)	5	(4)			(3)	3				(2)	(9)	322
United Kingdom	29	77	(4)	2	3	(4)	(11)	(4)	(3)	(6)	3	(1)			(9)	8	76
France		(4)	9	3		2	5	2	1	2					2	(3)	22
Netherlands	5	2	3	5		1	2	1								(2)	18
Japan	(2)	3			11	(2)										2	13
Switzerland	5	(4)	2	1	(2)	8	2								1	(4)	11
Germany	(4)	(11)	5	2		2	8	3	2	3	(1)				3	(5)	10
Spain		(4)	2	1			3	2		1					1	(2)	7
Sweden		(3)	1				2									(2)	5
Italy	(3)	(6)	2				3	1		2					1	(2)	2
Australia	3	3					(1)				1				(1)		2
Belgium		(1)															1
Singapore													1				1
Greece														1			1
Canada															(3)	(2)	(4)
Finland																17	(4)
TOTAL	322	76	22	18	13	11	10	7	5	2	2	1	1	1	(2)	(4)	491


Attribution of Expected Tracking Error (b.p.)

Figure 11. Risk budgeting — by country

- (5) The sum of the variance terms in the attribution is 448 bps — the difference between this number and the total tracking error is 43 bps! The benefits of diversification are, in this example, in fact a cost. It must be questioned whether this portfolio construction from the country perspective is efficient. Risk has been taken in the off-diagonals — is there any commensurate expected return? The biggest cost of diversification is in the co-variance between the U.S.A. and the U.K. portfolios.

6.8.5 Risk budgeting by global industry group

- (1) Rather than analyse the portfolio by country, it can be done just as easily by global industry group using our methodology. The results of this analysis are shown in Table 6 and Figure 12.

Table 6. Risk budgeting — global industry group

	Health care	Energy	Financials	Information technology	Consumer discretionary	Industrials	Telecommunication services	Consumer staples	Materials	Utilities	TOTAL
Health care	86	31	(7)	6	(10)	2	2	(6)	(7)	(5)	92
Energy	31	98	(7)	8	(16)	(10)	4	(6)	(20)	(4)	79
Financials	(7)	(7)	46	(6)	17	7	(1)	5	9	8	70
Information technology	6	8	(6)	79	(12)		3	(2)	(5)	(3)	68
Consumer discretionary	(10)	(16)	17	(12)	55	2	(5)	7	13	5	56
Industrials	2	(10)	7		2	29		2	8	4	45
Telecommunication services	2	4	(1)	3	(5)		26	(2)			27
Consumer staples	(6)	(6)	5	(2)	7	2	(2)	17	3	2	20
Materials	(7)	(20)	9	(5)	13	8		3	14	1	18
Utilities	(5)	(4)	8	(3)	5	4		2	1	8	15
TOTAL	92	79	70	68	56	45	27	20	18	15	491

Attribution of Expected Tracking Error (b.p.)




Figure 12. Risk budgeting — global industry group

- (2) Once again the sum of the off-diagonals is positive (amounts to 31bps), so that there has been a risk cost from diversification — is the portfolio manager expecting a commensurate return?
- (3) Figure 12 shows the same results that are displayed in Table 6. Clearly, there is a lot of risk being taken in the interaction between the energy and healthcare sectors.

6.8.6 Risk budgeting by stock

- (1) It was observed in ¶6.8.4 that relatively large amounts of risk were being taken in the interaction between the stocks held in the U.S.A. and the U.K. Using our methodology, it is relatively easy to view the risk budgeting in this area to analyse the precise sources of risk.


Attribution of Expected Tracking Error (b.p.)

Figure 13. Risk budgeting — U.S.A. and U.K. stocks

- (2) Figure 13 shows the apportionment of risk between the active stock positions in the U.S.A. and the U.K. Since there are too many stocks to display in the heat map whilst still being able to read the letters, we have shown the 25 most important (either positive or negative) contributions to the off-diagonal contribution to the total tracking error of 29bps.
- (3) It can be seen from the heat map that there are relatively large contributions to the off-diagonal risk arising from the holdings in General Electric, Citigroup, Exxon Mobil and Microsoft in the U.S.A. against HSBC, BP, GlaxoSmithKline and Vodafone in the U.K. The positions in these stocks are adding to risk because of the risk model covariance between them. There are other financial stocks that could be held in both the U.K. and the U.S.A. For example, does the portfolio

Table 8. Risk budgeting — by country within Europe

	United Kingdom	France	Netherlands	Germany	Switzerland	Finland	Spain	Italy	Sweden	Belgium	Ireland	Greece	Denmark	Portugal	Austria	Norway	TOTAL
United Kingdom	538	(25)	13	(78)	(29)	55	(29)	(43)	(20)	(10)	(4)	(1)	(5)	(2)		(5)	353
France	(25)	59	22	34	11	(17)	15	13	9	5	3	1	2	1		1	136
Netherlands	13	22	32	10	7	(13)	7	4	4	2							89
Germany	(78)	34	10	50	15	(30)	19	21	13	6	3	2	3	2		2	72
Switzerland	(29)	11	7	15	52	(26)	5	5	6								49
Finland	55	(17)	(13)	(30)	(26)	117	(12)	(12)	(11)	(2)			(2)			(1)	44
Spain	(29)	15	7	19	5	(12)	10	9	5	2	1		1				37
Italy	(43)	13	4	21	5	(12)	9	12	6	3	1		1				23
Sweden	(20)	9	4	13	6	(11)	5	6	5	1							22
Belgium	(10)	5	2	6		(2)	2	3	1	1							10
Ireland	(4)	3		3			1	1									7
Greece	(1)	1		2													5
Denmark	(5)	2		3		(2)	1	1									5
Portugal	(2)	1		2													4
Austria																	1
Norway	(5)	1		2		(1)											
TOTAL	353	136	89	72	49	44	37	23	22	10	7	5	5	4	1		858

Table 7. Summary risk statistics

Tracking error	8.58%
Portfolio beta	0.92
Correlation	0.88

manager have that much stronger view of holding Citigroup or General Electric rather than American Express, Wells Fargo or JP Morgan; is HSBC that much more preferable to Lloyds TSB or Barclays? He must have a stronger view to compensate him for the increased risk.




Figure 14. Risk budgeting — by country within Europe

Table 9. Risk budgeting — by sector within Europe

	Energy	Financials	Health care	Telecommunication services	Consumer discretionary	Information technology	Industrials	Materials	Consumer staples	Utilities	TOTAL
Energy	437	(70)	66	(2)	(57)	17	(39)	(41)	(17)	(18)	277
Financials	(70)	144	(45)	(17)	55	(9)	35	27	12	16	151
Health care	66	(45)	136	3	(11)	(10)	(9)	(10)	(15)	(13)	93
Telecommunication services	(2)	(17)	3	102	(8)	8	(3)	1			84
Consumer discretionary	(57)	55	(11)	(8)	54	(18)	30	21	5	7	78
Information technology	17	(9)	(10)	8	(18)	88	(10)	(8)	2	2	62
Industrials	(39)	35	(9)	(3)	30	(10)	20	16	4	5	48
Materials	(41)	27	(10)	1	21	(8)	16	19	5	5	37
Consumer staples	(17)	12	(15)		5	2	4	5	15	4	15
Utilities	(18)	16	(13)		7	2	5	5	4	6	14
TOTAL	277	151	93	84	78	62	48	37	15	14	858


Attribution of Expected Tracking Error (b.p.)

Figure 15. Risk budgeting — by industry group within Europe

6.8.7 Risk budgeting within Europe

- (1) We can drill down into the portfolio in many ways. In this section we look at the European component of the portfolio, and analyse the risk budgeting relative to the European component of the benchmark. For the purposes of this example, we look at the risk budgeting by country and by industry, although, of course, many other ways are equally valid.
- (2) Table 7 shows the summary risk statistics of the portfolio.
- (3) The expected tracking error of the European sub-portfolio in isolation is 8.58% — hardly surprising given that the total portfolio has a tracking error of 4.91%.

6.8.8 Risk budgeting by country

- (1) Table 8 shows the apportionment of the 858bps by country. It reveals both the risk incurred within each country and, most importantly, the effect of the co-variance terms between countries.

- (2) The table of results can also be represented in a heat map so that the relative importance of the contributions to the tracking error can be easily seen. The results are shown in Figure 14.
- (3) Figure 14 shows very clearly that the major risk position is in the U.K. active portfolio — this contributes 538bps to the total tracking error of 858bps. However, the sum of the entries in the leading diagonal is 878bps.
- (4) The results for Finland are interesting — the pure Finnish risk (in this case Nokia) is 117bps. However, the total effect on the European portfolio is 44bps, as the holding in Nokia is diversified in the other active positions.

6.8.9 *Risk budgeting by industry group within Europe*

- (1) The final analysis of this example is to do the risk budgeting by European industry group. The results are shown in Table 9 and Figure 15.
- (2) The active portfolio positions in energy contribute most to the tracking error — in fact some 437 bps. However, it is very interesting to note that the sum of the leading diagonal amounts to 1021bps. The portfolio has been efficiently constructed and has been able to re-cycle 163bps of diversification risk, presumably into areas where there is a commensurate expectation of alpha.

7. CONCLUSIONS

7.1 We have laid out an integrated approach to risk measurement and management. Our emphasis had been on de-mystifying the numerical barriers that typically surround this topic, so that portfolio managers can incorporate risk management into their daily routine. In the past the modelling approach and the associated tools available have not made this topic accessible or meaningful to portfolio managers.

7.2 The model represented in this paper is based upon co-variance, i.e. we have only modelled each stock with respect to its first and second moments of a normal distribution. Higher moments (e.g. skewness and kurtosis) have been ignored, but deserve further research. Similarly, the maximum likelihood framework could be extended to incorporate time-varying risk exposures. These are important extensions, as they would enable the techniques to be extended to fixed income and derivative asset categories.

7.3 Risk measurement is a small, but necessary, part of the process, and, in our view, is emphasised too much by investment consultants and clients. Risk management, the process of budgeting and controlling risk, is far more important.

7.4 Risk is not a bad thing unless there is no quantifiable expectation of a commensurate return. In this case it is merely business risk!

7.5 By separating the measurement of risk from its interpretation, we are afforded much more flexibility in our analysis without sacrificing the quality of our risk forecasting (measurement).

7.6 The approach suggested in this paper provides a repeatable systematic and disciplined process to assess the risk structure in equity markets and how it changes over time.

7.7 Using the same framework, we are able to account for risk in a portfolio, and hence spend the risk budget efficiently. By aggregating risk from the ‘bottom-up’, we are able to account for the entire budget without recourse to balancing items attributable to co-variance. We model stock specific effects explicitly (and mathematically efficiently) rather than using a stock ‘residual’ as a catch-all for other unexplained sources of risk.

7.8 We hope, by using modern interactive visualisation techniques, that the topic of risk management will gain broader appeal and cease to be the preserve of actuaries, investment consultants and quantitative analysts.

ACKNOWLEDGEMENTS

The authors would like to express thanks to their immediate colleagues and their employer for making this paper possible. However, the most important thanks go to our clients who ultimately prompted all the ideas contained in this paper as a means of helping to solve some of the problems that confronted them. In addition, we are grateful to the three anonymous referees who provided useful insight into the previous drafts of this paper.

However, none of the above should be held responsible for the content and ideas behind the paper nor for any errors or omissions contained herein — they are the sole and exclusive responsibility of the authors.

REFERENCES

- DEMPSTER, A.P., LAIRD, N.M. & RUBIN, D.B. (1977). Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society*, **39B**, 1-38.
- EVERITT, B. (1974). *Cluster analysis*. London U.K., Heinemann Educational Books.
- MARKOWITZ, H. (1959). *Portfolio selection : diversification of investments*. New Haven U.S.A., Yale University Press.
- PARKER, V.R. (editor) (2000). *Managing hedge fund risk : from the seat of the practitioner — views from investors, counterparties, hedge funds and consultants*. London U.K., Risk Books.
- RAHL, L. (editor) (2002). *Risk budgeting : a new approach to investing*. London U.K., Risk Books.
- ROSS, S.A. (1976). The arbitrage theory of capital asset pricing. *Journal of Economic Theory*, **13**, 341-360.
- SHARPE, W.F. (1963). A simplified model for portfolio analysis. *Management Science*, **9**, 277-293.

- SHARPE, W.F. (2002). Budgeting and monitoring pension fund risk. *Financial Analysts Journal*, **58**:5, 74-86.
- SHUKLA, R.K. & TRZCINKA, C.A. (1990). Sequential tests of the arbitrage pricing theory: a comparison of principal components and maximum likelihood factors. *Journal of Finance*, **45**, 1541-1564.

APPENDIX

CLUSTER TECHNIQUES

A.1 *Cluster Analysis*

Cluster analysis is a multivariate statistical technique for grouping objects in a manner that will help in the interpretation of those objects. The groups are mutually exclusive, and chosen in such a way that the members of each group are similar to each other while members of different groups are dissimilar. There are two principal ways of performing cluster analysis. These are hierarchical clustering and partitioning (non-hierarchical).

A.2 *Hierarchical*

A.2.1 Hierarchical cluster analysis produces a classification that has an increasing number of nested classes. The hierarchy can be illustrated by a two-dimensional diagram known as tree diagram or dendrogram. Figure A.1 shows a dendrogram based on sample data.

A.2.2 A dendrogram shows how closely matched each stock is to every other stock. The measure of closeness could be the correlation between stocks, and this is shown as the height on the dendrogram (a different

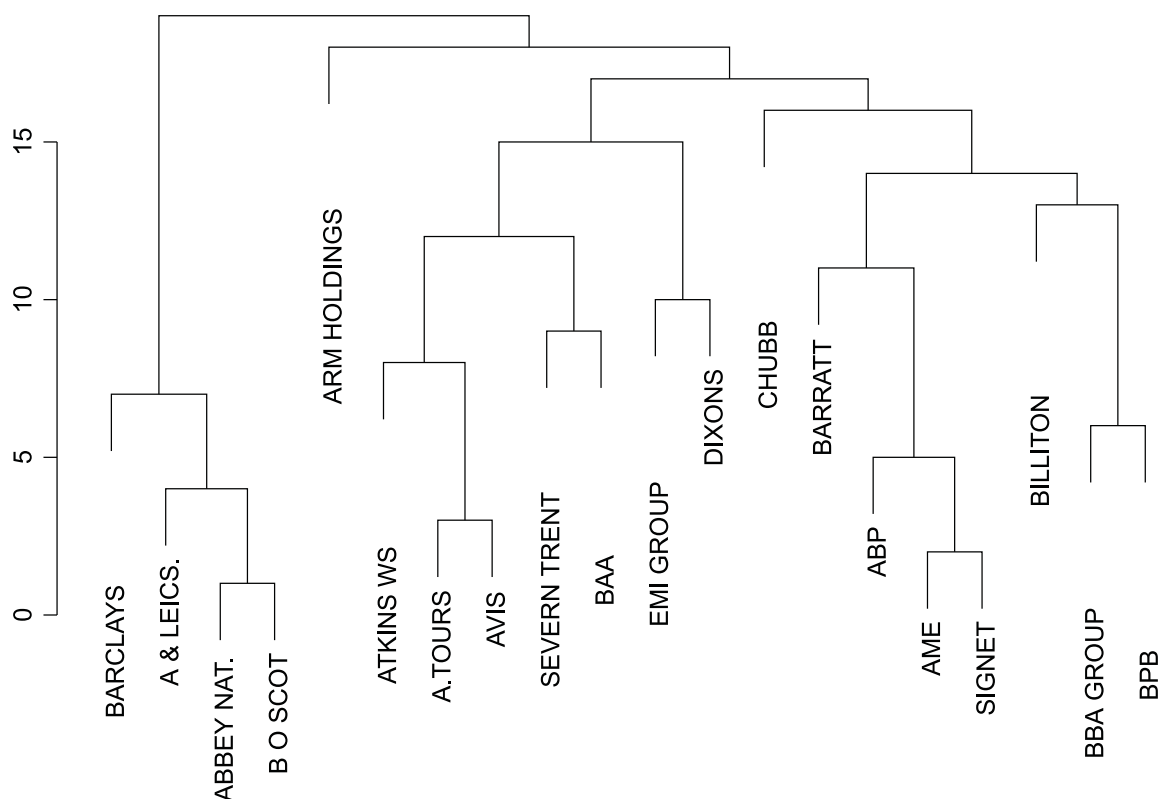


Figure A.1. Specimen dendrogram

measure of closeness has been used to produce the dendrogram above). In Figure A.1 the lowest intersection (between Abbey National and Bank of Scotland) shows that these are the closest stocks in the chosen universe. The next intersection up shows the next closest pairing, and so on.

A.2.3 Two different classes of algorithm — agglomerative and divisive — can produce dendrograms. Agglomerative algorithms operate by starting off with individual objects, then successively grouping them until all the objects are in one group. Divisive algorithms work in the opposite manner.

A.2.4 Once the relationship between all the objects has been established, we can specify a number of groups into which the objects should be classified. For example, if we required N groups, we would start at the top of the dendrogram and find the top $(N-1)$ intersections. The branches from these intersections would be the N groups.


A.2.5 The steps required in generalised hierarchical agglomerative clustering are:

- (1) identify what data are needed — e.g. stock returns;
- (2) compute the similarity/dissimilarity matrix (later in the appendix we look at different similarity/dissimilarity measures);
- (3) identify the two closest objects in the matrix;
- (4) merge the two closest objects into a cluster;
- (5) compute distances between the new cluster and all the other objects; and
- (6) finally, repeat the last three steps until all the objects are part of one cluster.

A.2.6 Different rules can be used to compute distances between two clusters. If the two closest members between the two clusters are compared, then the method is known as single link clustering. Several other methods exist — such as comparing the farthest members, comparing the group average and minimising the intra-cluster sum of squares (Ward's method). The sum of squares is used in the pan-European analysis.

A.3 *Partitioning*

Non-hierarchical cluster analysis requires a number of groups to be specified. Then it iteratively reallocates objects into groups until equilibrium is reached. No attempt is made to identify a link between the separate groups. For example, if we performed partitioning on a deck of cards and specified two groups, then we would expect to end up with one group of black cards and one group of red cards. Just as for hierarchical cluster analysis, there are several algorithms available for partitioning. The broad aim of these algorithms is to minimise the intra-cluster distance whilst maximising distances between different clusters. Figure A.2 shows the result of partitioning on sample data graphically.



These two components explain 100 % of the point variability.

Figure A.2. Partitioning on sample data

A.4 Other Important Considerations

A.4.1 As the objective is to group similar objects, a measure of similarity is required. For our analysis of stocks this could simply be the correlation between stocks. Stocks could be placed into an arbitrary number of groups by changing the threshold correlation values. However, correlation is quite crude as a distance measure, as it includes the specific risk component of stock return (as well as the market risk). In large, well-diversified portfolios we would expect specific risk of the stocks to be cancelled out, so a better distance measure would only consider the market (or non-diversifiable risk).

A.4.2 The distance matrix is constructed by comparing each combination of stocks. Our statistical factor model gives an exposure value to 20 orthogonal factors for each stock. We could compute the distance between two stocks as the Euclidean distance (sum of squares). Figure A.3 illustrates this in two-dimensional format.

A.4.3 We could use alternative definitions of distance, such as the angle between two objects — see Figure A.3.

A.4.4 The difference between this measure and Euclidean distance is that it allows for scalability. So, if object 2 were the same as object 1 multiplied by a scalar factor, then the second definition would produce a distance of zero, whereas the Euclidean distance would be non-zero.




Figure A.3. Difference between distance and angular measures

A.4.5 Initially we ran our analysis using both measures, and found that the Euclidean distance measure was better at segmenting the data into meaningful groups. If we think about the problem in factor analysis terms, then this result is not surprising. If exposures of object 2 were twice those of object 1, then we would expect the two to behave differently in the context of equity returns.