Consistent Assumptions for Multinational Asset Models

Jeroen van Bezooyen
Keith Guthrie
Robert Howie
Peter Ludvik
Shyam Mehta
Andrew Smith

Abstract

Building an international ALM requires assumptions about expected returns on bonds, equities and currencies around the world. How can we be sure these are consistent?

We test several methods of deriving assumptions. Using simple methods, we discover that the larger the model gets, the greater the chance that even sensible looking assumptions imply major market inefficiencies.

The question is, can we find a better assumption set that passes the efficiency test but where the numbers still make some sort of intuitive sense? Should we bodge the mean returns, or the variance covariance matrix - or do we have do invent some other psychological bias to explain standard assumptions?

1. Introduction

Stochastic asset models are continually becoming more comprehensive in the asset classes covered. Initial models of equities and gilts have been extended to include cash, index linked bonds, foreign exchange rates, real estate and whole yield curves. Much current research focuses on consistent modelling of various derivatives, and on the yield spreads associated with various levels of liquidity and credit risk.

Earlier models had a modest number of parameters; the requirements for consistency have become well understood. But as the number of parameters grow, it is less clear that judgment alone is adequate protection against hidden inconsistencies.

Let us take a simple example of choosing correlation matrices. For illustration, let us think of a correlation expert, who selects correlations randomly between –1 and 1. A correlation matrix must be positive semi-definite, but what of these random matrices? The risk of the expert selecting inconsistent correlations grows dramatically as the matrix size increases. For even moderately sized matrices, say 10 by 10, the odds are overwhelmingly against a random matrix being positive definite.

In the same way, are there hidden pitfalls in the construction of multinational asset models?

2. Purpose of Multinational Modelling

2.1 Overview

Many actuarial consulting firms, fund managers, investment banks and investment consultants have built or use a multinational stochastic asset model. Sometimes this will be a relatively simple model (with some specification of expected mean returns together with a variance covariance matrix); occasionally rather more complex asset return, price and yield dynamics are modelled. Such models are used for a number of purposes:

- Comparing risks and executed returns for different multinational portfolio choices, typically with the ultimate goal of
- Choosing an 'optimal' portfolio mix, often having regard to the current portfolio distribution and to prior constraints. For example, one might wish to test the effect on risk and return of hedging currency risk, of increasing the overseas content in a portfolio, or of matching to the fixed cash flow profile of the domestic liabilities
- Risk management, for example to identify hedging strategies and to quantity risk
- For some applications, the model will be used as a pricing tool, for example to assess the cost of a multi-index basket option.

More recently, Exley, Mehta and Smith (1997) have suggested that alternative strategies can be compared having regard to a value maximisation principle: the model is used to quantify and value 'second order' transaction cost and other effects in order to determine the strategy which provides the maximum value to the client. The concepts here are (1) that the client can himself buy or sell in the primary securities markets and hence swapping one asset for another through a typical ALM study adds no value, and (2) that clients prefer more to less.

2.2 Efficient Frontier Analysis ('EFA')

It will be evident from the above discussion that the primary use of multinational models hitherto has been to in one form or another construct an efficient frontier, or risk/return trade-off framework, in order to assist with portfolio choice. Although the need for portfolio choice decisions is clear, the lack of robustness of EFA is perhaps less well known amongst the actuarial community.

- A small change in the expected rate of return on an asset class is typically associated with a very large change in the optimal asset allocation, and
- Estimates of expected rates of return vary widely from person to person.
  In this environment EFA and similar approaches will not assist in asset allocation decisions. However, rather than reject EFA many consulting firms vary the assumption set until the desired allocation appears optimal.

2.3 Individual versus Institutional Asset Allocation

In Exley, Mehta & Smith (1997) it was suggested that institutions (life companies and pension funds, for example) could profitably adopt market value based decision making rules, as noted in 2.1 above. EFA does have a role, but largely in respect of portfolio choice for individuals rather than institutions.

An individual choosing, for example, his overseas equity allocation will need to consider several aspects:

- Consistent assessment of the expected returns for the various asset classes. How robust are these estimates?
- Marginal contribution to risk to his portfolio from different allocation possibilities
- Own portfolio definition: is the risk analysis to include his own property, mortgage, labour income wealth, …..? How do overseas currencies and returns correlate with his income, for example?
- His risk aversion for different levels of his (real?) income and wealth in different possible future status of the world
- Liquidity and other implications of investing overseas. Is he at an informational or tax disadvantage compared with local investors?

3. Current Approaches to Setting Assumptions

3.1 Setting Expected Returns

As noted above, the expected return decision will typically be crucial in driving the results of any analysis. A difference of ¼% p.a. expected return differential between the UK and US equity markets, or between bonds and equities, will have a major result on the analysis.

Practitioners propose various ad-hoc techniques, for example

- All equity markets have the same return
- Returns equal to the averages over some prior period
- Returns set equal to those estimated by an equity market strategist or economist
- Expected excess returns are proportioned to local equity market volatility or,
- Expected returns are set according to some international CAPM benchmark.
- More complex models such as APT (arbitrage pricing theory) may be used

Closer analysis shows that none of the proposed techniques stack up for the intended purpose: the fact is that a wide variety of expected return assumptions are perfectly plausible. International CAPM, for example, may be a reasonable benchmark but is simply not accurate enough for EFA purposes.

3.2 Setting Correlations and Volatilities

It is often believed that setting the covariance matrix is a relatively robust process, at least compared with the impact of the expected return assumption on EFA results. However, there are a number of complications:

- Perhaps most importantly, there is increasing awareness of extreme risk events as being a major factor in portfolio choice. An 'average' equity market standard deviation of 20%pa may in fact be made of 'normal' volatility of 10%pa combined with an occasional crash. Indeed in a paper on Market Efficiency, EMS (1996) estimated that some 50% of the equity risk premium could be attributable to crash rather than 'normal' market volatility.
- Stress events are typically associated with very different correlation patterns than in more normal market environments.
- More generally, there may be a variety of underlying factors with differing distributions (not necessarily normal) driving international returns. Constructing and using the covariance matrix of total returns is inefficient if there are better relationships to consider. For example, if purchasing power parity applies in the long-term, a model which did not reflect this could be in danger of overestimating exchange rate risk and in mis-specifying hedge ratios. A model of nominal interest would similarly be inefficient if there are important links between inflation rates and real interest rates.

3.3 Risk Neutral Assumption Sets and Home Country Bias

The difficulties referred to in 3.2 can be overcome if a pure market value driven decision process is in place. Excess returns are offset by the corresponding market risk costs, and the starting point for every return assumption is the risk-free rate for the home currency. Given that investors appear to have a 'home country' bias, it may be appropriate to layer on top of this risk-free rate a negative convenience yield effect to reflect the additional compliance costs and loss of customer appeal from investing overseas. Clearly the amount of any convenience yield adjustment will be very arbitrary, even more so if there is an attempt to modify the adjustment according to local market characteristics (country rating, size of market, distance from the home market...).

Notwithstanding the logic of <u>not</u> using an EFA approach, except when providing advice to individuals, (who will typically be unable to afford the advice!), in the rest of this paper it has been assumed that a 'real world' or 'EFA' type analysis is required.

4 Random Walk Models

4.1 Risk – Return Models

In this paper we concentrate on the simplest possible models: multivariate lognormal random walks. To avoid inconsistencies with maturing fixed income instruments, our multivariate random walk is applied to the ratio of total return indices. One of these total return indices must be selected as the numeraire, or accounting unit. Other indices are then expressed relative to the numeraire.

We assume that the log of these total return indices is a multivariate random walk, whose variance-covariance matrix is *V* per unit time. We suppose that the drift of the random walk is ì - ½diag(*V*), where diag(*V*) denotes the vector whose elements are taken from the diagonal elements of *V*. We choose this notation so that the (arithmetic mean) expected returns on each element are given by the elements of exp(ì)-1.

Unlike the situation with correlation matrices, any possible choice of ì could be used as the basis of a projection model. To the extent that there are constraints, these emerge as difficulties in interpreting model output. For example, consider portfolio optimisation by utility maximisation subject to positive holding constraints. For any given utility function, we usually find that the optimal portfolio usually omits some asset classes. But suppose that, after trying a variety of different utility functions and investment structures, some asset classes are unattractive to everybody. Then we would start to question the model assumptions. We might suggest that the expected return assumptions for the unpopular assets had been selected too low.

These difficulties often only arise a long way down the road of model use. Model assumptions with embedded unpopular assets may look reasonable to the naked eye. The unpopular assets would not be picked up in statistical analysis of simulation output. They would not be noticed in the first model application, nor perhaps in the second or the third. Indeed, to be fully confident that unpopular assets were not hidden inside a model, requires extensive optimisation testing.

4.2 Numerical Examples

To illustrate the ideas, we took looked at total return indices from MSCI. Our data consisted of monthly returns from the seven years 1994 through 2000.

The series examined were:

Cash indices	Bond total return	Equity total return
	Japan Germany UK US	Australia Canada France Germany Hong Kong Italy Japan Netherlands Singapore Spain Sweden Switzerland UK US

We hope to repeat these investigations with more extensive series.

We measured returns in excess of the US bond total return index. We took the historic volatilities and correlations as given, and focused on possible choices of mean returns. We investigated six different ways of choosing the mean returns:

- Based on seven years' history
- Zero risk premium for bonds, fixed risk premiums for equities in all countries
- Risk premium proportional to volatility
- CAPM, based on US equities as the market portfolio
- APT, setting all bond risk premiums to zero
- APT, assuming equal risk premiums for all European markets

The resulting risk premiums (arithmetic mean basis) are shown in the chart below:

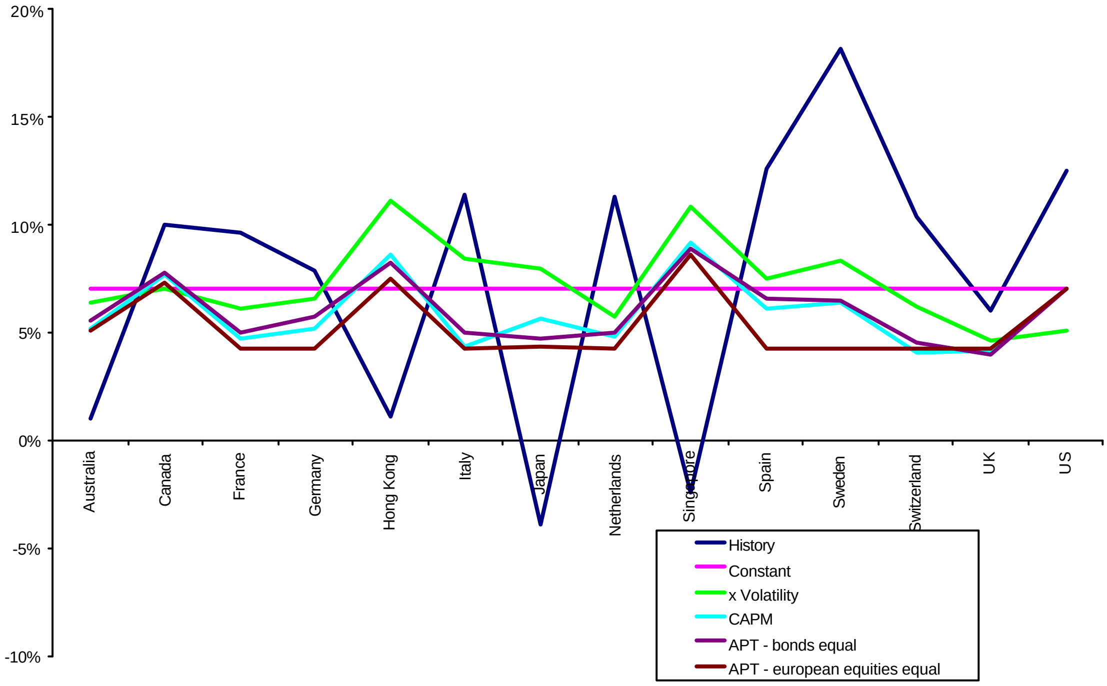

We can see some themes here. Efficient markets reward systematic risk only, which is why the theoretical models do not necessarily attribute richer risk premiums to volatile markets. More volatile economies often carry less systematic risk because much of the volatility is diversifiable. Notice the similarities between CAPM and APT results – this is not surprising as APT is a generalisation of CAPM.

4.3 Sharpe Ratios

One way of comparing the efficiency of economic models is to examine optimised Sharpe ratios. See the accompanying technical note on market efficiency for more details of this process. The chart below shows the Sharpe ratios

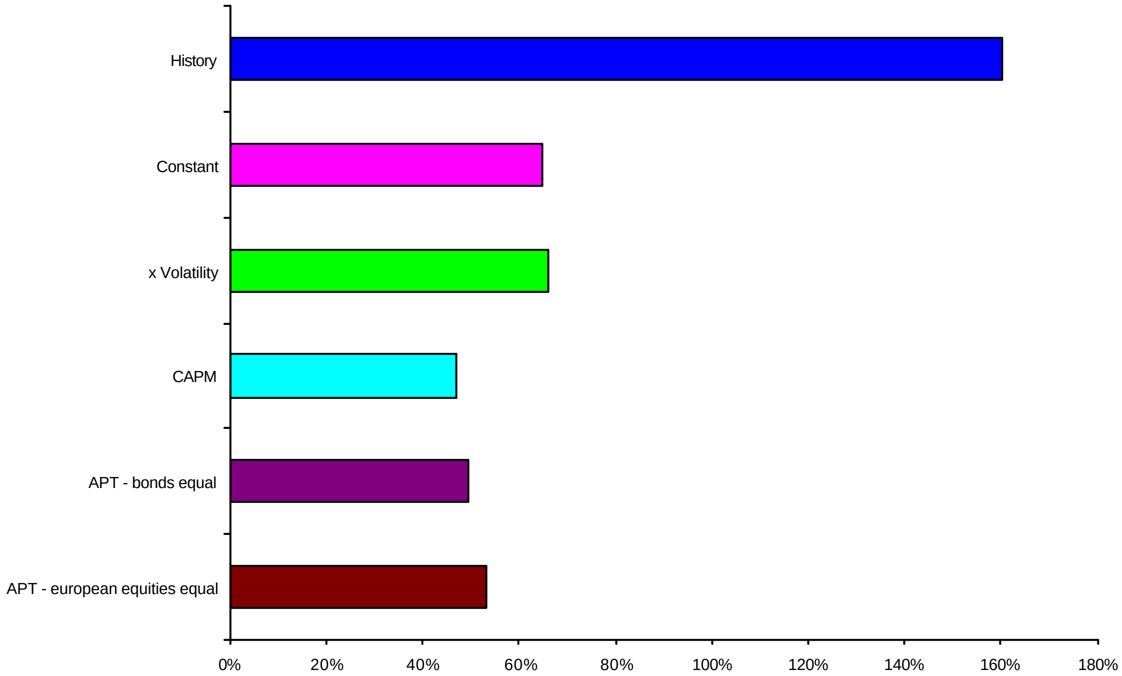

We can see here that the use of historic means to estimate the future is likely to give the least efficient market. This is not a proof of market efficiency – on the contrary, this highlights the amount of noise in the raw data if no return model is applied. Using volatility weightings actually produces a *less* efficient market than assumptions of constant risk premiums by asset class. The model approaches generate more efficient markets than the ad-hoc rules.

4.4 Compromise Solutions

We have looked at ways of choosing optimally efficient models. Our methods involved picking assumptions that are as efficient as possible, subject to some imposed input constraints.

A more general approach is to start with a set of risk premiums, perhaps based on historic data, and then to nudge them in the direction of a more efficient combination. Mathematically, we are seeking to optimise some combination of low Sharpe ratio and closeness to a set of data-driven assumptions.

4.5 More General Models

The lognormal random walk in this paper is possibly the simplest model of asset prices. Generalisations would allow for

- conditional risk premiums varying over time
- conditional volatilities and correlations varying over time
- non-normal distributions with skewness, kurtosis and so on

Our methods can be applied to these more general cases, but some care is needed. The relevant means and variances should be the cross-sectional, conditional moments. We cannot ensure an efficient market by forcing relationships between long term returns on asset classes. Efficiency implies fair pricing at any point of the economic cycle. That is why conditional moments are important.

In practice, these conditional moments may not be easily accessible mathematically. It is also difficult to estimate conditional distributions from classically generated Monte Carlo simulations. As a result, it is difficult to apply our techniques to models such as the Wilkie model. Instead, our ideas suggest structures that might be put in place when choosing a new model.