Revised and retitled June 2007

Multi-scale variation, path risk and long term portfolio management

Roger J. Bowden and Jennifer Zhu*

Abstract

Strategic fund management and similar investment styles may be exposed over time to multi-scale variation that can create inconsistencies between portfolio methodology and welfare objectives. This paper utilises a more embracing temporal frame of reference that can reconcile macro-scale variation in investment value with the shorter run, according to welfare preferences as to path risk versus terminal value. Path exposures give rise to path risk, and this is operationally defined in terms of a spectral welfare function, encompassing Fourier and wavelet analysis within a common framework. The approach leads to dynamic analogues of mean-variance such as band pass portfolios that are more sensitive to variability at designated scales.

Key words: Band-pass portfolios, path dependence, spectral utility functions, strategic portfolio management, value at risk, wavelets.

JEL classifications: G11, G14-15, G23; E32, E44; C32

Bowden is with the School of Economics and Finance, Victoria University of Wellington, P.O. Box 600, Wellington, New Zealand. Phone + 64 4 472 4984, fax + 64 4 472 4983, emails: roger.bowden@vuw.ac.nz; roger.bowden@kiwicap.co.nz . Zhu is with the Dept. of Finance, Banking and Property, Massey University at Wellington Private Box 756, Wellington, New Zealand. Phone: +64 4 801 5799 email j.zhu@massey.ac.nz. Comments from two *Quantitative Finance* referees have been of material assistance in focussing ideas and exposition, and are gratefully acknowledged. Thanks go also to Finance seminar participants at Massey, Otago, Waikato and Auckland Universities.

1. Introduction

The objective of this paper is to establish an explicitly multi-scale approach to portfolio methodology that enables a better reconciliation between measurement and design methodology, on the one hand, and underlying welfare objectives, on the other. The methods that result are well adapted to asset accumulation in an economic environment that creates longer run, or macro scale, variation on which is superimposed shorter run fluctuation arising from market noise and similar disturbances. Investment outcomes are compared in terms of the value paths they create as distinct from a more narrow focus on either short run or longer run returns in isolation. Path risk merges as a secondary welfare objective, arising from path exposures and measurable in the form of a penalty functional of spectral power, in the present application a function of wavelet energies at different scales. What results can be viewed as a dynamic analogue of mean-variance analysis, though applied to investment value as such, rather than to any pre-selected return definition or measurement. Band pass portfolios can be designed to optimise the primary long run welfare function subject to user preferences as to path risk. The development has points of contact with a number of other areas of financial management, such as path dependence in derivatives pricing, and value at risk.

Distinctions between terminal value and path risk come to the fore in the context of strategic fund management, a style that typically emphasises longer period returns rather than shorter[1]. It could be argued that strategic portfolios should be geared to longer run macroeconomic or structural influences, rather than the shorter run dynamics of financial markets. Doing so, however, entails some methodological issues and agendas.

(a) Are returns as conventionally measured the most effective basis for portfolio selection? Even in a shorter run context, fractionally integrated representations for long memory processes (Granger (1980), Baillie (1996), Bhardwaj and Swanson (2006)) suggest the use of fractional differences of log values as a more appropriate return concept. Things become more complicated when the focus shifts to long-period value growth. Use of five or ten-year compound returns is certainly possible,

[1] Most investments textbooks have a treatment of this and other management styles, as do investment advisers specialising in the area such as Ibbotson Associates, drawing on early papers such as Ibbotson and Sinquefield (1976). Recent technical discussions include Campbell and Viceira (2002) and Guidolin and Timmermann (2004). Strategic management can rebalance portfolios from time to time in response to investor circumstances or structural economic changes, which we do not consider in the present contribution.

but the implicit geometric average smoothing is not necessarily optimal, even if it could be decided which is the most appropriate holding period return to use. It might be better in all cases to resort to investment value or log value as the starting point, and devise methodology that allows a more flexible return concept to emerge naturally from the data or from the economic context.

(b) Measurement and portfolio methodology should be mutually integrated, as two sides of the same coin. The body of techniques used to measure value growth should also be usable for purposes of portfolio construction. If the measurement techniques are well adapted to the longer run movements associated with business cycles or economic growth, the same output should be used to construct portfolios that can exploit longer run sources of investment value.

(c) Strategic portfolios commonly focus on maximisation of the expected value (or expected utility) of long run portfolio value; or equivalently, long run average returns over a stated horizon. But objective functions of this type do not say very much about what happens along the way. Two alternative portfolios might have just the same terminal or long run mean-variance properties but yet be exposed quite differently to alternative path dynamics, e.g. to the sequence of one period conditional distributions along the way. In terms of the primary objective function this may not matter, but there may be implicit secondary considerations under which it does. Thus paths that exhibit excessive long swings in portfolio value could be exposed to investor withdrawal, which is damaging to the reputation of its managers. Portfolio selection needs to address issues of path exposures and path risk, including the relationship with constructs such as value at risk.

(d) Methodologies adapted to the above exigencies might also be suitable for other agendas. If we can design portfolios suited to long run strategic investment, the logical complement is a portfolio that captures the short run; both short and long run portfolios then emerge as special cases of a more embracing technology. Likewise, we could think of designing portfolios that are suited to different informational habitats on the part of investors. Some managers might have better economic information or models, and hence prefer the longer run end of the spectrum; others might feel more at home with shorter run financial market dynamics.

Wavelet analysis offers a convenient tool for the above purposes. In its measurement aspect, it represents a decomposition of a given series into bands of successively longer run variation, generalising such stationary techniques as complex frequency demodulation in spectral analysis. Though the core ideas have been around for many years, wavelet theory and practice experienced a renewal of interest and development through the work of authors such as Daubechies (1988), (1990), (1992), Mallat (1989), Coifman *et al* (1990), and Cohen *et al* (1992). For useful reviews of the use of wavelet analysis in economics, see Ramsay (1999, 2002), Schleicher (2002), or Crowley (2005). Wavelets have also found application in finance, e.g. Capobianco (2004), Lee (2004). Cognate applications are Gencay and Selcuk (2004), who note that high long term volatility is not necessarily accompanied by high short run volatility, while In and Kim (2006) suggest the use of different hedge ratios according to time scale in risk management.

In the normative aspect developed in the present paper, one can use the wavelet energy decomposition to design portfolios tailored to preferences as between long or short run variation. The techniques could be regarded as generalising mean- variance to a multi-period context, in which a subsidiary preference structure is established by weighting differentially the spectrum of short to long run variation, generalising the idea of spectral utility functions (Bowden (1977), Otrok (2001)). Band pass portfolios can be designed to exclude altogether designated long or short-term value fluctuations.

The scheme of the paper is as follows. Section 2 introduces some of the ideas that motivate subsequent development, notably path risk and the issue of how to measure changes in portfolio value in a way that is consistent both with history and the requirements of portfolio optimisation. Section 3 opens with a very brief review of wavelet ideas and the way that they will be applied. The asset classes to be used are introduced and their wavelet properties summarised. Section 4 contains the core normative development of the paper. Band-pass and similar portfolios are obtained by a process of analogy and comparison with classic mean-variance. Issues such as the choice of primary objective function are explored. Discussion concludes with an illustrative comparison of the result with mean-variance. Section 5 contains some concluding remarks.

2. Motivation: path risk, and the empirics of investment value

Multi scale behaviour over time in investment behaviour is typically associated with the existence of shorter run influences such as market noise, superimposed on longer run macroeconomic determinants such as those associated with business cycles, monetary and fiscal policy, or exchange rates[2]. The first part of this section explores some of the problems this creates for long term portfolio management such as that associated with strategic modes or related passive style portfolio modes. The notion of path risk provides a convenient framework to do so. Also considered are the operational aspects of the units of measurement (values or returns), with some further introductory remarks about wavelet analysis as a framework for the measurement of path risk.

2.1 Path risk

Path risk arises as a consequence of path dependence, which is reflected in the time inseparability of utility functions. Thus in a fund management context, the primary objective function might be cast in terms of a terminal distribution of unit value, but the system dynamics entails a sequence of conditional distributions generating an interim dynamics that can entail incidental or secondary welfare effects. The wedge between primary and secondary welfare effects could arise for a variety of reasons. It might be a consequence of dual responsibility, as in the distinction between investors and management, each with specific objectives in addition to those shared. Alternatively, the distribution functions of the primary objective might be easier to estimate, perhaps because they are asymptotic in nature and consistent with a variety of different formulations for the sequential conditional distributions, which may themselves be difficult to specify and estimate with any exactness. In such a case portfolio design should allow for possible adverse effects from alternative path histories. The following remarks are intended to elaborate on the general idea and define the metrics to be used.

Let $(\Omega, \Im, P)$ denote a probability space, and $\{\Im_t; 0 \le t \le T\}$ a filtration representing the information stream available to the manager. There is an $\Im_t$ - adapted control process $x_t(\omega)$ that can be taken to represent the portfolio choice at any point in time, and denote by $X_t = \{x_u; u \le t\}$ the history of control settings up to time point t . A static or passive portfolio policy is a special case in which x_t is determined at the outset

[2] Unhedged exposure to exchange rates, in particular, can be a potent source of multi-scaling behaviour. For a wavelet study see Bowden and Zhu (2007b).

and is constant thereafter[3]. There is an adapted process V such that $V_t = \nu(X_t, t, \omega)$, called the 'unit value history', that will be taken to represent the accumulating portfolio value per dollar of initial capital invested. For simplicity, all proceeds are assumed continually reinvested, i.e. no external dividends until terminal time T. Finally, there is a utility function U defined on the alternative path histories, which can remain general at this point.

Time inseparability is taken to mean that at any intermediate time $0 < \tau < T$, conditional expected utility from that point on is a function not only of V_τ, but of the entire history of unit values to that point. In other words, the dynamic programming principle does not apply. Extension of the state vector is a possible recourse. For instance, the managerial utility function could reflect in addition the number of investors in the fund, for if investors exit as a consequence of poor unit value performance, this will impact adversely on managerial income or even employment. This would require explicit modelling of investor responses to unit value histories, which is not something that the typical fund manger would want to attempt. A less formal approach is to recognise that some types of path history are going to be more exposed to investor unhappiness and exit (in this example). Such paths - essentially portfolio choices x_t - might then attract a penalty, and their reward element must correspondingly compensate. So as with any form of risk analysis, there are two aspects, first to recognise the path exposures, and second to provide a metric that will penalise paths that are more exposed.

Spectral analysis provides a possible framework of this kind, interpreted broadly to cover Fourier analysis and wavelet decompositions. A constructive analogy with classic mean-variance analysis is useful. In a long term mean-variance framework, expected terminal utility could be written as $U(\mu_T, \sigma_T^2); U_1 > 0, U_2 < 0$, i.e. where μ_T and σ_T^2 are respectively the mean and variance of terminal unit value. Now time domain path variability is expressible in terms of the Cramèr representation (stationary Fourier analysis) or the scaling detail decompositions in the case of wavelet analysis. Correspondingly, let E stands for a spectral density ordinate for Fourier analysis, or

[3] For a formal distinction between static and dynamic portfolio modes, see Haugh and Lo (2001), Kohn and Papazoglu-Statescu (2006). In their treatment, static portfolios are essentially 'fire and forget', and in a complete market are equivalent to some derivative position established at the outset. Dynamic portfolios are those that need continual rebalancing in response to the particular historical path. The present application is implicitly one of market incompleteness rather than completeness, but the ideas are analogous.

detail energy for wavelet decompositions, and suppose that G is the set of available energy ordinates for the application in hand. The required path variance analogue would then be a positive definite function φ: G$\rightarrow$ $_{+}$. In addition to overall energy, the function φ is chosen to assign penalty weightings to the frequency or detail elements that are designed to accord with managerial perceptions of secondary exposures or some other type of path exposure. In the case of wavelet decompositions, the mean element can be translated into the highest order approximation, i.e. that which remains once most of the details have been removed, completing the analogy with mean-variance. Portfolio efficiency in this extended framework requires maximising the expected reward element, subject to acceptable values of the path risk metric φ.

In what follows, path risk will be defined operationally as a weighted sum of spectral power ordinates, of the form

$$\varphi = \sum_k w_k E_k ; \quad w_k \geq 0 , \sum_k w_k = 1 . \tag{1}$$

The weights $\{w_k\}$ are set by the user in accordance with perceptions of path exposures. The following discussion illustrates some of the possible considerations in choosing whether to penalise at the longer or shorter run end of the power spectrum; specific implementation is contained in section 4.

Path preferences could be couched in something as simple as a general preference for smoothness. Paths that soar can also plunge, and the welfare benefits and costs might not compensate: investors are more apt to anxiety on extended downturns than euphoria on the upturns. As to the choice of weighting function, much would evidently depend on the advertised stance of the fund. High long term energy (variation) would be more a danger to a fund that advertised itself in terms of stable balanced growth.

A more structured approach to path welfare costs might run in terms of a dynamic version of value at risk, or in this context, value at risk duration. The idea is that paths should not spend too long below a comparator or benchmark path that could be interpreted as a moving VaR critical value, incorporating investor regret. Figure 1a depicts path histories generated by two alternative static portfolios, both valued in terms of US dollars. Portfolio A is an accumulation index for NZ stocks[4], while B is made up of an Australian equity index portfolio and a US treasury bond index in equal

[4] For data definitions see section 3 and Appendix D.

proportions. In addition, path C is a benchmark portfolio representing investment in US treasury stock held to maturity, using the 10 year yield at inception time (6.01% at May 1993) as a proxy for the entire period 14 year rate. All unit value series are measured in logs, hence the straight line for the benchmark. The two subject paths A and B are not widely different as to the variation of monthly returns. Note, however, that path B is uniformly above the benchmark C, but path A spent almost three years below (see the double headed exposure duration arrow). Unit holders might tolerate short dips, but not long baths. Even though path A is a little superior on a primary welfare objective taken as terminal unit value, its adverse performance on the penalty element would likely see managers prefer path B. In terms of ex ante portfolio selection, portfolios such as A, viewed as more subject to long swings, would carry the burden of higher path risk, for which additional compensation would be required in the trade-off with the primary welfare objective. Moreover, a point such as P has different welfare consequences for path A versus path B, and in a dynamic portfolio problem could lead to different portfolio choices from then on, as a manifestation of utility inseparability. Further points of connection are with path dependence in options pricing, and with the notion of path dependence in economics and political science, the idea that 'history matters'.

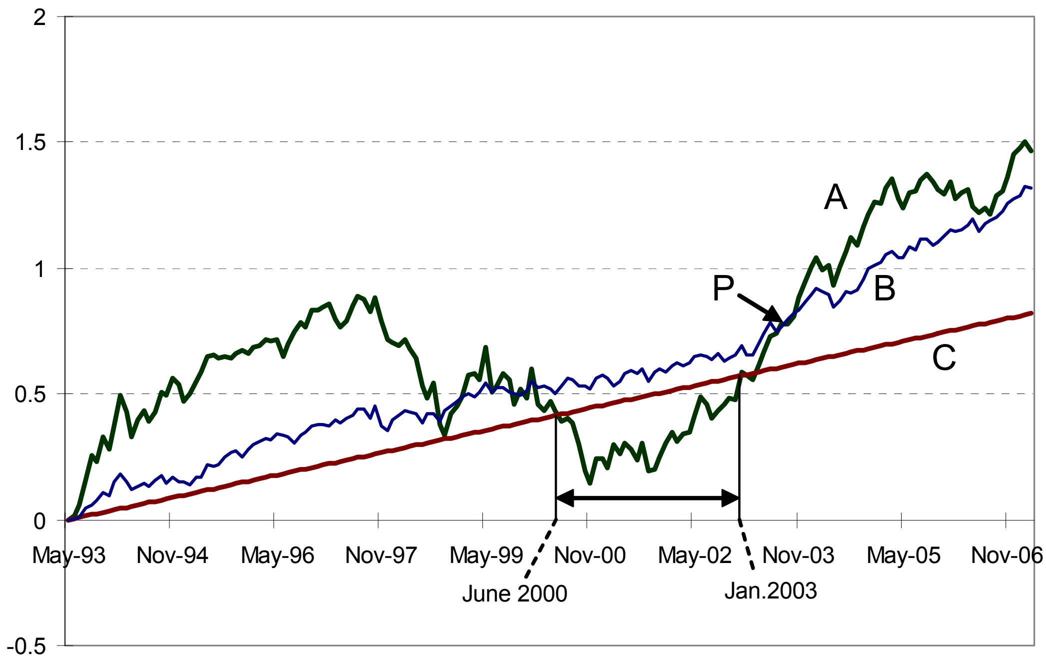

Figure 1a: Path risk from benchmark violations

The role of the weighting function (1) is illustrated more explicitly in figure 1b, which reallocates the total detail energy of path A towards the finer of the 7 detail levels (wavelet conventions and methodology are exposited in section 3.1 below). Total detail energy of path A is 11.984, of which the four bands spanning 2 years and longer contribute 11.574 for the original series. Series B is a reconstruction of A that redistributes the total detail energy so that only 4.830 remains at the longer fluctuations. The result is evident in the form of greater short run variation. Notice, however, that the duration of the adverse exposures of path B (below path C) are transitory, never lasting longer than two to three months at a time. Given the implied preference as to reward and path risk, we could say that history B is path-preferred to A. Of course, it might not be so for an alternative choice of the weighting function in expression (1) for path risk. Thus short run fluctuations might well be viewed as less attractive for this particular class of fund. But that is a matter of managerial choice.

In the remainder of the paper it is shown how the weighting function approach can be built into a portfolio selection procedure. In order to add anything to just the total variance, and hence just path independence, the detail energies have to exhibit clear scaling behaviour. The extent to which this is true in the given investment context is explored in section 3.

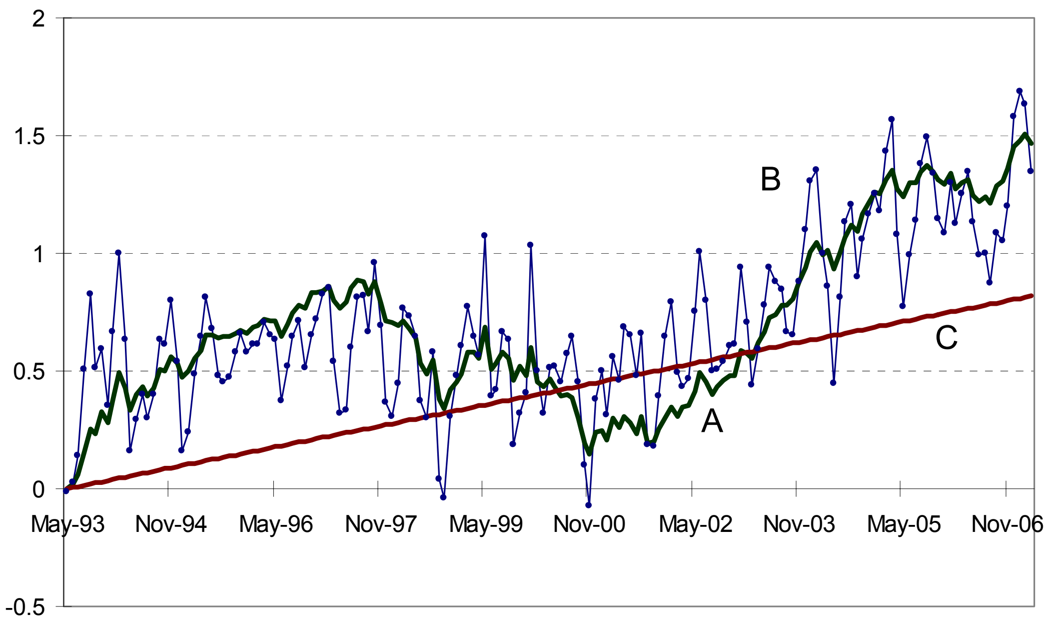

Figure 1b: Energy structure and benchmark exposure duration

2.2. Values versus returns

Portfolio analysis has traditionally been formulated and conducted in terms of security returns, for historical and other reasons. To facilitate subsequent discussion, assume that the discrete time interval is small so that returns can be treated as synonymous with changes in the log of value. Recent developments in stochastic modelling have raised an issue as to whether security modelling and portfolio design should not utilise alternative constructs. As an instance, suppose security values are generated in terms of a fractionally integrated time series of order *1+d*, where *d* is the fractional component $0<d<1$. Such representations have been used in as a way of capturing long memory time series (Candelon and Gil-Alana (2004)). A more suitable return concept in such a case might be based on $(1-L)^{1+d} \log V_t$ rather the conventional $r_t = (1-L) \log V_t$, where *L* is the backward lag operator.

Alternatively, one might like to preserve an open mind about any specific return definition and work at the outset in terms of $\log V_t$ itself or some scaling transformation such as $\sqrt{1/t}\log V_t$, which can be useful where geometric Ito processes are thought to generate the data[5]. Logically, there is no difference in using returns or log values; they contain exactly the same information. But inference based exclusively on returns can suffer from a swamping with shorter run market noise, making it harder to accept as statistically significant macroeconomic effects that work more slowly but are important over years, rather than weeks or months. Short-interval rates of return are typically dominated by market transients, including not only one-period white noise, but transients such as local bubbles or market overshooting in response to news or sentiment. Transient effects of this kind have been described by authors such as Barnett *et al* (1989), Lux (1995), Kojima (2000), Scheinkman and Xiong (2003). Market transients can generate episodic short term serial correlation that will magnify the variance of short term returns, but become smoothed out over longer horizons, allowing underlying fundamentals-based signals to become manifest. Thus wavelet analyses of the rates of return typically show energy increasing with detail – the higher the detail, the greater the energy. The effect shows up in the present data (see Appendix B). A similar effect has been noted in the literature on efficiency losses or gains from the use of overlapping observations, where the implicit smoothing enables better detection of

---

[5] Suppose that returns are generated by the continuous time geometric process $dV_t = \mu V_t + \sigma V_t dB_t$ where B_t is a standard Brownian motion process. As $Var(B_t) = t$, it follows that $\sqrt{1/t}\log V_t$ can be decomposed into a trend element of order $\sqrt{t}$ and a zero mean detail $\sigma B_t / \sqrt{t}$, which has constant variance (energy, in the wavelet context). Thus one indication that the log value series is behaving like the classic geometric accumulation process is that once normalised by $\sqrt{t}$, the details should show no obvious expansion in amplitude over time.

longer run movements (e.g. Fama and French (1988), Campbell and Shiller (1988), Boudoukh and Richardson (1993)). The situation is analogous to business cycle indicators, which are often prepared in two forms, the levels version and the growth or 'growth cycle' version, the latter referring to percentage rates of change. Locating cycles and their turning points is easy with the first, but quite difficult with the second.

2.3 Using wavelets

Wavelet analysis offers a flexible approach to some of the issues raised in the preceding, without requiring stationarity in returns, or too much regularity in the cyclical behaviour of investment value.

(a) Wavelet energy decompositions are an excellent way to represent empirical path risk. Paths with higher energies in longer-term details (cycles) might be thought more likely to trigger investor exits or other adverse management exposures. By an appropriate choice of portfolios one can redistribute a given total energy over the frequency spectrum, accepting higher short run fluctuations as the price of diminishing the kind of exposures that matter, namely to longer run swings. Or on the other hand, investors might prefer quite a different kind of history profile where they can bear longer swings, but are averse to shorter run loss of value.

(b) Wavelet algorithms embrace both levels and returns in a single framework, without needing to sharply distinguish them, or to choose one as against the other. The nature of the filter used at each detail is effectively a graduation as between differencing and accumulating. For instance, the level 1 detail in the Haar filter is essentially a first difference: applied to log values, it would amount to the returns as conventionally defined, while progressively lower levels of detail ultimately end in the approximations, which correspond to a trend in values. Filters such as the Coif and Sym waveforms can be viewed as a more flexible approach to differencing and smoothing operations. It is no longer necessary to make a formal distinction between trends, stochastic or otherwise, and fluctuations, or to establish very precisely the order of integration or cointegration of the given data series.

3. Wavelet decompositions of asset classes

The main objectives of this section are to define the asset classes to be used in the portfolio analysis (section 4) and to illustrate their wavelet properties. Wavelet analysis itself is now fairly well established in economics and finance (Ramsey (1999, 2002) and Crowley (2005)), so only a brief review follows. Readers unfamiliar with the techniques may like to consult Percival and Walden (2000) and Gencay and Whitcher (2002) for more extended expositions.

3.1 Wavelet decompositions

A wavelet is rather like a sinusoid localised at a particular point in time, so that its amplitude drops off rapidly on either side of that time point. A choice of wavelet forms is available, depending on which seem to conform more closely to the typical cycle of the data in hand, or to the typical notion of a business cycle as analogous to a sinusoid of irregular periodicity and amplitude. In the present study the Coif 5 wavelet was used as the default choice; the Sym 10 waveform was also employed, with much the same results. Moving along through time, one fits a succession of such wavelets. Each time point contains contributions from wavelets of the same 'scale' (quasi frequency) but centred at neighbouring points. In addition, it will also contain contributions from wavelets of different scales, corresponding to cycles of different frequencies. By a similar mathematical argument to complex demodulation, one can express the series at any point in time as a sum of the wavelets of different scales. The shorter scales represent higher frequency fluctuations, while the large-scale wavelets capture the long run movements. Collectively across different scales, the wavelets of either family are flexible enough to allow for asymmetric local cycles of rather arbitrary form, so this is no longer a story requiring regular sinusoidal patterns as in standard Fourier analysis.

Although chosen from the same generic family, the wavelets of different scales are normalised to refer either to the cycles ('mother wavelets') or long term trend or quasi trends ('father wavelets'). The results of fitting mother (cyclical) wavelets of different scales are called the 'details' (D) and they are additive in their effect. Progressive sums, by adding more details, are called the 'approximations' (A). Level 1 is the smallest scale or highest quasi frequency, so D_1 represents the cycle at this highest level of detail[6]. The given series is then split into $D_1 + A_1$, where A_1 is the series once the very shortest fluctuations have been removed. Levels 2,3…. contain successively less small-scale complexity. Extracting these leads to broader time frame approximations designed to reveal longer run cycles and ultimately the trend.

[6] The reader should be aware that some treatments adopt the opposite convention so that detail 1 would then be the longest scale.

An 'average period' construct for a given level of detail D can be derived by finding the sinusoid whose period most closely matches that of the wavelet fitted at any point in time, suitably adjusted for its scale. Then one simply averages out these local equivalent periods over time. This enables us to think of the successive details as corresponding to progressively longer cycles, just as in spectral analysis. Appendix A illustrates the average period convention used in the present study.

As with spectral analysis, one can measure the amplitudes of each cyclical component in the form of the variance of the wavelet detail. Unlike spectral analysis this is a local concept, differing over time. However, one can use compute the average variance over the given time horizon and present the results in the form of a table of average wavelet energies (AWE) at the different levels of detail.

3.2 Some energy decompositions

The top row of table 1 gives the asset classes used to construct the illustrative portfolios. The vantage point is that of an international equity portfolio for a New Zealand investor wishing to invest in the US and other major stock markets. The total return stock indexes as given refer to the own-currency log value, e.g. the US index is in US dollar terms, the Japanese component in yen.

The remaining asset appearing in table 1 is the total return index on a one-month forward contract on the US dollar against the NZ dollar. This corresponds to a portfolio long in zero coupon NZ bonds or bills[7], short in US, embodying the foreign exchange (FX) hedging component of the three-fund theorem of international finance (Solnik (1974)). The foreign exchange hedge component assumes particular significance when the stock returns are to be converted back to home currency (here the NZD) as noted above, as indeed they have to be in the present case. In that case the bond portfolio can be referred to as a currency hedge portfolio against the USD, and is in effect a forward contract.

The relevant monthly return on the hedge asset is defined by

$$r-(1+r^*)\frac{e_1-e_0}{e_0}-r^*,$$

where:

r = NZ one-month bank bill rate as of start of month

[7] The one month NZ bank bill rate (BBR) is chosen, while the US short rate is taken as the certificate of deposit (CD) rate. Both are wholesale rates referring to high-grade bank credit.

r^* = one-month US CD rate as of start of month

e = exchange rate as 1USD = e NZD, i.e. with the US dollar as commodity currency and NZ dollar as terms: e_1 = end of period rate, e_0 = beginning of period.

The hedging monthly returns are then accumulated to form a total return index homologous with those used for the equity components of the portfolio. The object variable (or dependent variable) in each case is the log of the total return index.

Table 1 gives AWE decompositions for the asset classes used in the present study using monthly data from Jan 1988 to May 2007 (see Appendix D for data definitions and sources). The Coif 5 wavelet is used, and computations were done in Matlab (Misiti *et al* 2005) using Mallat's algorithm. Very similar results were obtained with the Sym10 filter. The maximum detail available for the data run is level 7. What is left over is taken to be the trend, though it may remain more complex than the standard log linear trend. Indeed this is one of the strengths of wavelet analysis, that it makes no pre-judgements about the form of any underlying deterministic or stochastic trend.

Table 1: Average wavelet energy decomposition for asset classes

Detail Level	Equivalent time period (months)	NZ stocks detail energy (%)	US stocks detail energy (%)	Japanese stocks detail energy	Australian stocks detail energy (%)	USD/NZD forward detail energy (%)
1	2.9	2.29%	1.09%	1.88%	2.50%	0.58%
2	5.8	2.75%	1.30%	2.28%	2.91%	0.96%
3	11.6	4.51%	2.02%	4.31%	7.06%	1.06%
4	23.2	6.48%	2.81%	12.20%	12.14%	2.19%
5	46.4	4.94%	6.30%	20.43%	17.54%	8.44%
6	92.8	57.55%	57.51%	54.63%	38.53%	43.11%
7	185.6	21.49%	28.97%	4.27%	19.32%	43.66%

The power pattern in most cases shows an interior peak at level 6, which corresponds to an average cyclic period of 8.5 years. Note the higher energy of US stocks at longer run energies, especially detail 6. By way of contrast, applying the Coif and Sym filters to the standard exponential Ito model with constant drift reveals an absence of interior local energy peaks. Power accumulates monotonically through different detail levels, effectively becoming a stochastic trend (Appendix C illustrates).

3.3 Time series representations

Time series plots of the details are given in figures 2a,b for two of the assets, namely NZ stocks and the forward rate hedge asset. The log value cycles are irregular both as to amplitude and period.

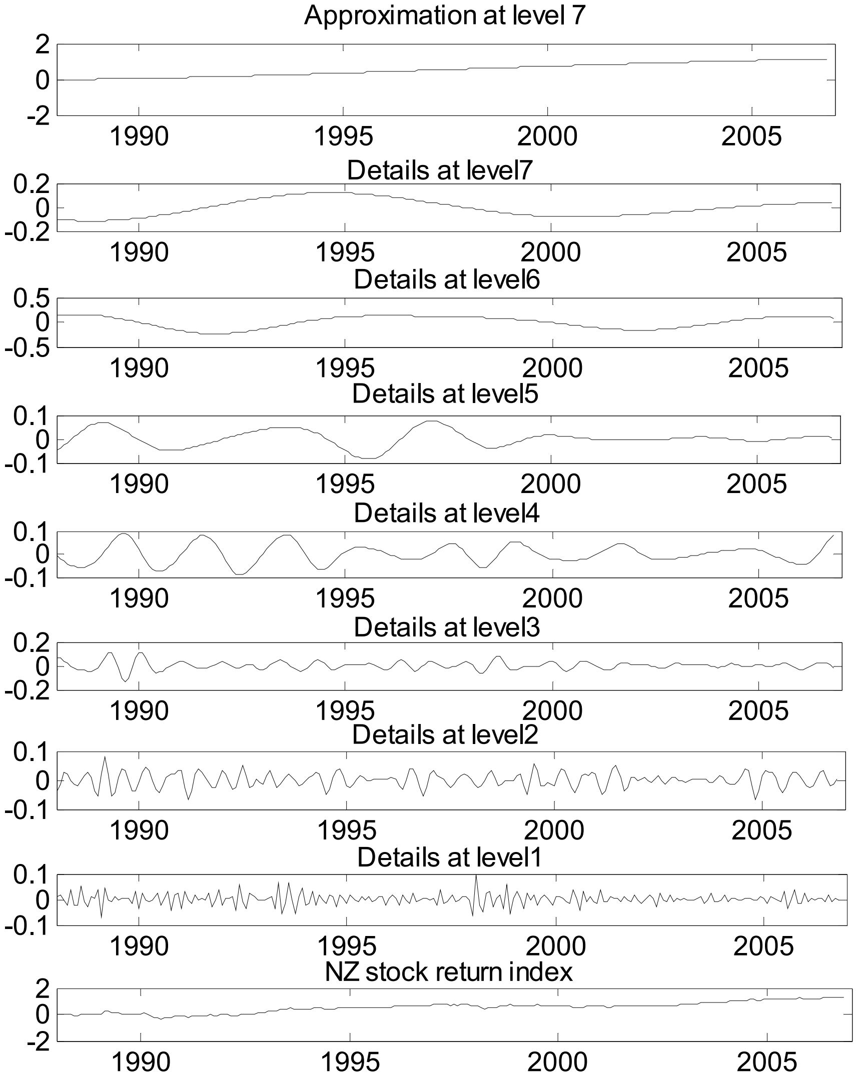

Figure 2a: Wavelet decomposition of NZ stock index accumulated return

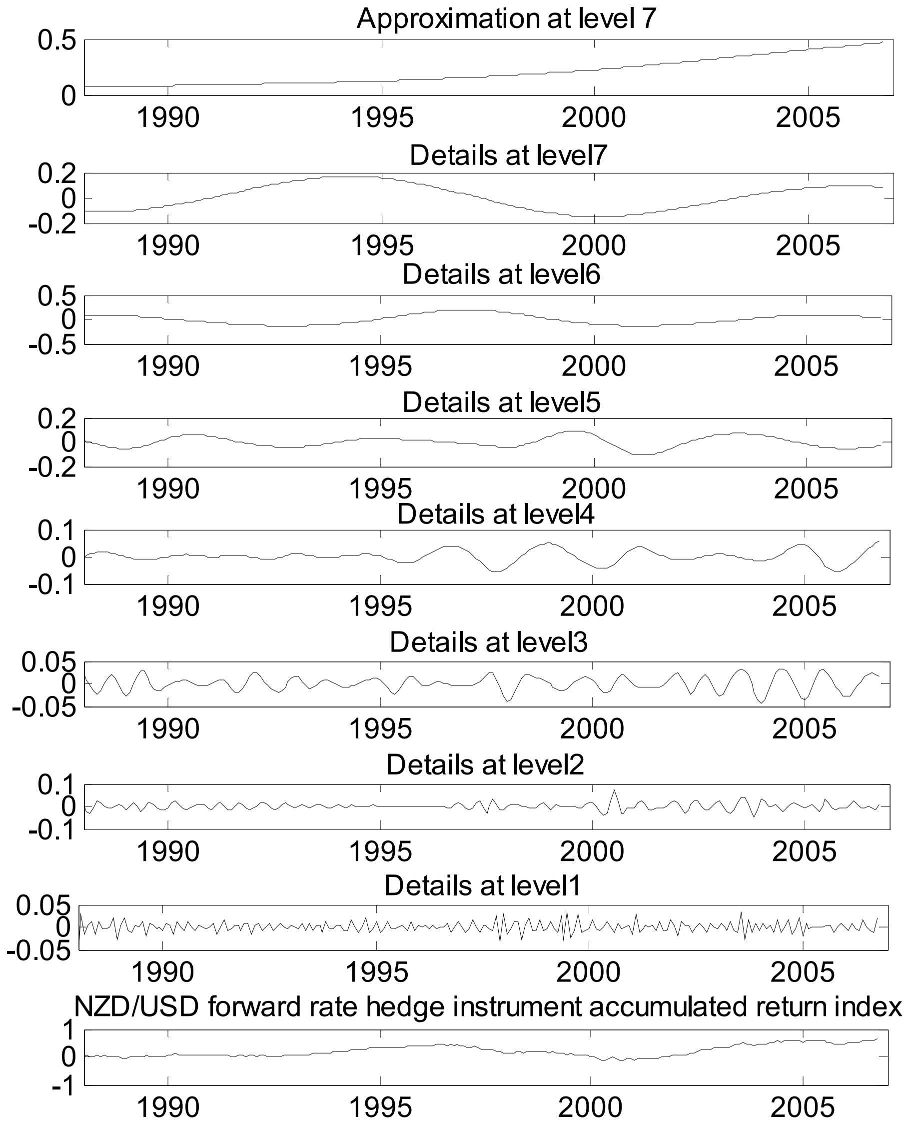

Figure 2b: Wavelet decomposition of USD/ NZD forward rate hedge instrument accumulated return

4. Portfolio choice

This section develops the band pass and related portfolio ideas, implementing the ideas discussed in section 2 and illustrating with the asset classes of section 3. Generalising the familiar mean-variance portfolio construction, wavelet-based portfolios maximise long run reward subject to limits on the energy penalty (φ). The collection of such points is the 'reward-energy frontier', analogous to the efficient frontier of mean variance.

By combining assets into a portfolio, one aims to achieve paths that are less risky for any given level of reward. In the wavelet based approach, we choose to carry out the portfolio analysis in terms of (log) values directly, and not one-period returns as such. However, there is a natural relationship between the two. If a set of assets {i} have values V_i, then the asset proportions are $x_i = V_i / V$ and log accumulation per period is taken as $\sum_i x_i \Delta \log V_i$ which is a portfolio return. In some formulations, return elements can appear in the objective function (see below). However, in what follows wavelet approximations (A) and details (D) refer to log portfolio value.

4.1. Reward objective and path risk penalty

In a context of wavelet analysis, the reward objective for long run investment is most naturally taken as some metric assigned to the high level approximations, such that higher values at any time point are preferred to lower values. Let A_* denote the wavelet approximation of maximal order consistent with historical data availability, taken as T observations. For brevity, we shall sometimes loosely refer to this as the trend. In the preceding sections this was taken as A_7, as in figures 2a,b. A suitable objective might then be of the form

$$\text{Max } T\sum_{\tau=1}^{T} f(\tau)\Delta A_{*\tau} \quad , \tag{2}$$

which is a weighted sum of the historical value increments, with $\Delta A_{*\tau} = A_{*\tau} - A_{*\tau-1}$.

In expression (2), $f(\tau)$ is a non-negative weighting function such that $\sum_{\tau=1}^{T} f(\tau) = 1$. It can be chosen to accord with some preference as between early or late return accumulation. For instance, a preference for earlier returns could be represented in the form of a discount factor. Some other useful special cases are as follows:

(i) The uniform weighting function $f(\tau) = 1/T$ all τ, corresponds to the usual geometric rate of return over the whole horizon. Multiplying this by T as in the above objective gives $A_{*T} - A_{*0}$. Hence the objective is simply to maximise the terminal value of the trend. If the wavelet decomposition is carried out on logs to begin with, then the objective is the compounded value growth. If the factor T was missing in the objective (2), and logs were used, then the objective would be interpreted as the long run trend geometric average rate of return.

(ii) Take $f(\tau) \propto \theta^{T-\tau}$; $0 < \theta < 1$. The higher the value chosen for θ, the more one weights later values of value growth. The idea behind this is that later value increments of the historical record might have more predictive content for what is to come in the present real time.

The path risk penalty function φ is assumed to be linear, with nonnegative coefficients $\{w_k\}$, so the constraint is of the form:

$$\varphi = \sum_k w_k E_k \leq v\,; \quad w_k \geq 0\,, \sum_k w_k = 1\,. \tag{3}$$

In expression (3), E_k denotes the average energy at detail level k, and ν is a user-assigned constant, interpreted as a path risk constraint. Thus by setting some of the weights $\{w_k\}$ to zero, and assigning heavy penalties to the others, the resulting portfolios will favour variation in the former, but not the latter. Once could call these band pass portfolios, motivated by similar usage in electronic system design, where one filters out signals at designated frequencies, allowing others to pass through unhindered. The empirical illustration below adopts the quasi VaR stance outlined in section 2.1, where the fund manager is influenced by a need to avoid by investor fear of large-scale opportunity losses associated with long swings in value. In this case the longer details are penalised, and the weightings assigned to short run variation are small or even zero. However, the general band pass approach can handle quite different weightings. For example, a fund manager concerned that excessive short run fluctuations in a volatile market might unsettle investors, could elect $w_k = 0$ for $k > 2$. This would allow lower level details to pass through unhindered while penalising short run fluctuations $(k = 1,2)$.

4.2. Optimisation problem and equivalent utility function

The optimisation problem is to choose the asset weights to maximise (2) subject to (3). Also incorporated are standard portfolio constraints, such as asset proportions have to add up to unity if they require capital, or be non-negative if fund policy requires this. By varying the allowable energy parameter ν, and solving the resulting portfolio, one can trace out an efficient frontier in just the same way as for classical mean-variance analysis.

An equivalent utility function is

$$U = T\sum_{\tau=1}^{T} f(\tau)\Delta A_{*,\tau} - \lambda\sum_{k} w_k E_k \ , \tag{4}$$

where in the programming context, $\lambda \geq 0$ is interpreted as a Lagrange multiplier. As noted in section 2.1, there is a parallel with the mean-variance utility function, in this case via expression (4). Using the second mean value theorem, the effect of a weighted sum of energies is as though there is a single energy E_*, say, which in turn has the dimension of a variance. So the equivalence relationship with mean-variance (denoted ~) can be expressed as:

$$\varphi = \sum_{k} w_k E_k = E_* \sim \sigma_T^2 \tag{5a}$$

$$T\sum_{\tau=1}^{T} f(\tau)\Delta A_{*,\tau} \sim \mu_T \ . \tag{5b}$$

The equivalence of the weighted sum of energies with a variance is useful in choosing the allowable path risk constant ν in the programming specifications (see below).

Efficient frontiers can be obtained by varying the allowable path risk parameter ν and finding the portfolio that maximise the reward function (2) subject to the weighted energy constraint (3). Plotting the reward against ν, or loosely just φ, yields the efficient frontier incorporating the trade-off between reward and path risk. It will be referred to in what follows as the reward-energy efficient frontier.

4.3. Application

The portfolios are constructed with the asset classes appearing in table 1 section 3. These are intended to be operational portfolios, so each of the stock returns have been converted to home currency. The portfolios are constrained to non-negative proportions to the four country stock market portfolios, and there is a single zero capital element, namely the USD/NZD forward contract, which can be shorted, i.e. potentially have a negative portfolio weight. In the objective function we used the time weight $\theta = 0.9$ which has a mean distributed lag of 9 months, indicating that most of the value increment weight is assigned to the last 18 months of the horizon. For the energy weights $\{w_k\}$ we assumed equal weights for details 4-7 but zero weights for energy levels 1-3, i.e. that investors are unconcerned about short run fluctuations. By way of contrast, in standard mean- variance analysis, investors are assumed to be equally worried by short and long run power elements. In the present band pass application, we are allowing power bands 1-3, i.e. the shorter run fluctuations, to pass freely.

Figure 3 depicts the resulting reward-energy frontier. The efficient frontier is strikingly similar to that of standard mean variance analysis, with the same parabolic shape extending into the lower inefficient half. As with mean variance, the trade-off (implied λ value) is higher as the energy bounds diminish.

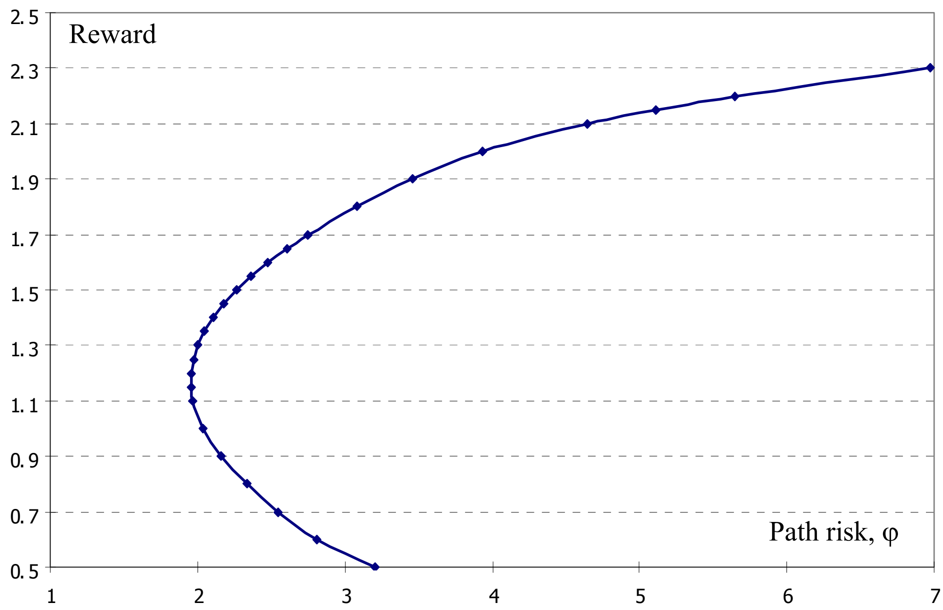

Figure 3: Reward-energy efficient frontier

Table 2 gives the optimal asset weights as one moves along the efficient frontier. Note that these do not have to add up to unity because of the presence of a zero capital element, namely the forward contract. Only the stock weights add up to unity. As the energy bounds become more restrictive, the optimal portfolio rebalances to Australian stocks (with lower long scale volatility), while the proportion devoted to US stocks decreases. It is also evident that the use of the USD/NZD forward contract diminishes. At first sight, this looks puzzling; one might have expected a forward to be variance diminishing. However it is consistent with the finding of Thorp (2005), also Bowden and Zhu (2006), that preserving an exposure to unhedged spot USD/NZD is actually conservative risk management practice in Australasia.

It is of interest to see whether the similarity with mean-variance extends to the path properties of the optimal portfolios. If the two give similar results, this could be taken as comfort in the use of mean-variance analysis. If not, then the issue of optimality and long-term stability would have to be addressed, with the REF portfolio as a useful starting point.

Comparison of the two approaches can never be rigorous, as they refer to different reward or variation concepts, and standardisation is needed on one or the other. Thus if we assume that locally dependent returns may be possible, then mean-variance assumptions of return independence are not satisfied, so that wavelet analysis is the more appropriate. However, returns can certainly be computed for a mean-variance analysis, so there is nothing to prevent an operational comparison of the two. For the wavelet approach, the specifications are chosen as above for the reward and energy weightings θ and $\{w_k\}$. Two alternative comparisons along these lines are as follows.

(a) Normalise on the MV version of the reward i.e. a given mean return over the entire horizon. Select the corresponding portfolio along the REF frontier that generates this mean. Compare the time paths and energy decompositions, of the two portfolios, including the path risk, defined for this purpose as the sum of the four longest detail energies as in expressions (4) or (5a). Figure (4a) illustrates, while table (2a) gives the energies, and the respective portfolio compositions.

(b) Normalise on the path risk φ. Start with a MV efficient portfolio and calculate its path risk. Find the portfolio along the REF frontier that has the same path risk and compare its mean and reward with that of the original MV portfolio. Figure (4b) plots the two time paths of the resulting portfolios, while table (2b) contains the energies and portfolio compositions.

A third possible approach (not illustrated here) might be to normalise on the trade-off parameter λ between reward and variation, with appropriate interpretation of these dimensions in the respective contexts. Variation would be taken as σ^2 for MV and as weighted energy for the REF portfolio.

Normalising on the mean as in (a) shows the smoothing effect of the wavelet based approach, reflected in the lower value for the path risk metric φ in table 2a. The REF portfolio was slower to rise between 1996-2000, but with much less of a subsequent fall. The effect is apparent in the lower detail 6 energy and in the relative path risk (2.73 as against 3.68). In portfolio terms, it is produced by down-weighting the US stocks component in favour of Australian stocks. The latter have virtually the same mean as the US, but materially lower long-term variation. The hedge proportion allocated to the US dollar has also diminished. The Japan weighting disappears altogether.

Normalisation (b), with fixed path risk, results in a higher accumulation path for the REF portfolio. As before the Australian weight is increased at the expense of US stocks, but the tendency to use USD/NZD forwards remains roughly the same.

Table 2a: Mean-normalised comparison between MV and REF portfolios

	Mean-variance efficient portfolio MV	**Reward-energy efficient portfolio REF**
Assets		
NZ	15.96%	18.81%
US	54.78%	19.62%
JP	9.90%	0.00%
AU	19.35%	61.57%
Forward	57.98%	20.22%
Reward	1.5966	1.6854
Path risk φ	3.6806	2.7333
Total detail energy	3.8985	3.0544
Single Period Mean	0.0098	0.0098
Single Period Variance	0.001183	0.0015

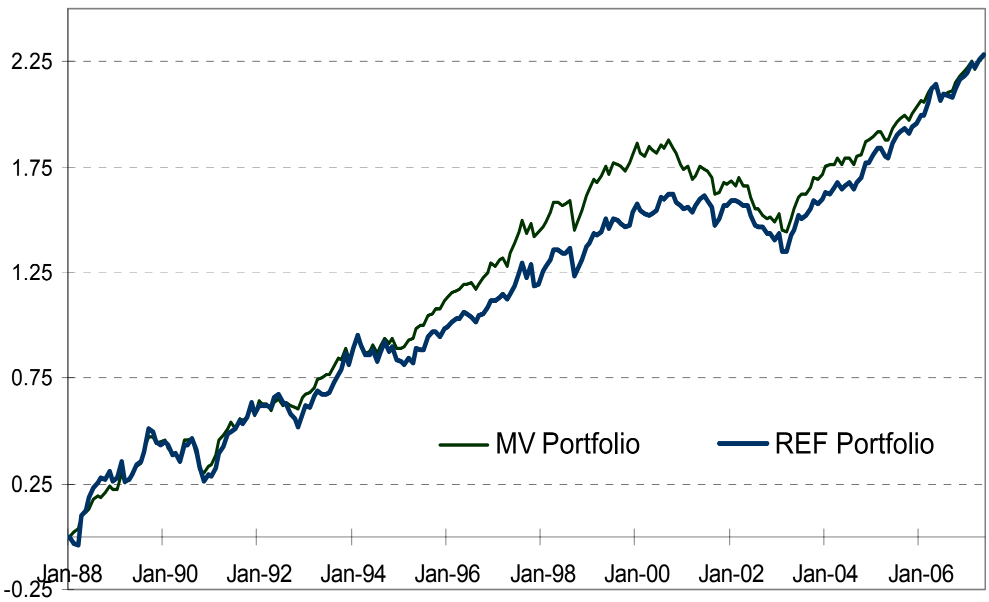

Figure 4a: Path comparison, MV and REF portfolios: mean-normalised

Table 2b: Energy-normalised comparison between MV and REF portfolios

	Mean-variance efficient portfolio MV	**Reward-energy efficient portfolio REF**
Assets		
NZ	15.96%	0.00%
US	54.78%	16.76%
JP	9.90%	0.00%
AU	19.35%	83.24%
Forward	57.98%	44.79%
Reward	1.5966	1.9511
Path risk φ	3.6806	3.6806
Total detail energy	3.8985	4.0462
Single Period Mean	0.0098	0.0115
Single Period Variance	0.001183	0.0017

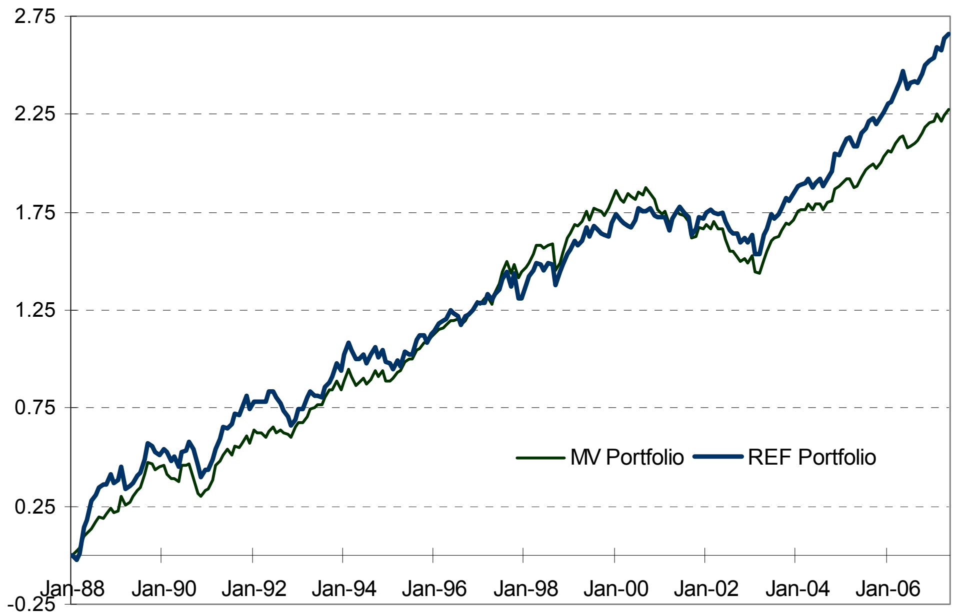

Figure 4b: Path comparison, MV and REF portfolios: energy-normalised

4.4. Extensions

Preceding development assumes that the manager's objective is to maximise long-term reward while minimising path risk. A quite different scenario might be that of a hedge fund concerned with identifying portfolios that actually maximise path risk over some designated detail band, perhaps one much shorter than assumed above. This could be accommodated by requiring a minimal reward element – or even deleting it – and maximising the path risk with an appropriate choice of the energy weights $\{w_k\}$. This looks a bit like a dual formulation from the classic theory of mathematical programming (Dantzig 1963, Rockafellar 1968, Murty 1976). However, the latter would require one to minimise the energy subject to reward constraints, so the 'hedge fund problem' is not quite dual to the long run strategic approach.

5. Concluding remarks

The underlying objective of the paper has been to develop portfolio technology that allows for non-independent return elements, dependence that may be weak in the short run but have a cumulative impact over the long run. There are other ways of attempting the same thing, notably by developing formal models of conditional returns, but they do require additional macroeconomic or time series modelling, which may be difficult. Explicit structural or time series modelling also requires the manager to vest a lot of faith in the predictive performance of the model, which is problematic, given that economists are not all that good at long range forecasting of business cycles. The wavelet based reward-energy approach is much less demanding in assumptions or informational requirements than formal causal modelling. It can be viewed as generalising mean variance analysis to the spectral domain, where the latter is interpreted broadly to encompass possible nonstationarity and wavelet technology. The energies correspond collectively to the variance in mean-variance analysis, but refer to path properties as a whole rather then the property of a single stationary distribution.

On the other hand, the wavelet based approach does have some maintained hypotheses of its own, notably that the long-term volatility patterns are characteristic of the data generation process, e.g. an underlying business cycle, and are likely to be repeated in the years to come. There is some comfort in the ability of wavelet analysis to detect structural breaks, which typically appear as sudden energy bursts in the high detail bands (Vuorenmaa (2005), Bowden and Zhu (2006)). On the other hand, Ramsay (2002) has noted that wavelet decompositions may not be stable outside the sample period. That is always going to be a potential problem, for the use of higher frequency data does not necessarily help to identify long cycles. However, it may be possible to endow the wavelet approach with *a priori* macroeconomic information, with the objective of increasing confidence in the long-term volatility patterns. For instance, if asset returns and value accumulation depend on exchange rates, one might have a fair idea about the causal factors involved for a particular home country. In terms of our example, the NZ dollar is well known to be driven by world commodity price cycles in conjunction with monetary policy responses (Lance (1996), Bowden and Zhu (2007a,b)). Likewise, there is some support among NZ economists for a business cycle of about 7-8 years, partly as a result of the commodity-exchange rate cycle (e.g. Kim *et al* (1995), Hall and McDermott (2006)).

Further points of contact in finance are with factor models of returns. In standard factor models such as the APT model, asset returns are generated in terms of unobservable orthogonal factors, but the latter continue to require the efficient market accumulation model. In the wavelet-based approach, the factors are orthogonal but no longer necessarily temporally independent. As with APT, there may exist dual portfolios that embody the factors. It would be of interest to explore what such portfolios might look like in different capital markets, and how they could be exploited in funds management or even as a basis for generic classes of fund.

Appendix A

Wavelet scale and frequency

The purpose of the discussion that follows is to establish a convention for measuring the equivalent average wavelet length of a detail, which is useful for tabular presentations.

To connect the wavelet scale to frequency, a pseudo frequency is calculated. The algorithm works by associating with the wavelet function a purely periodic signal of frequency F_c that maximizes the Fourier transform of the wavelet modulus. When the wavelet is dilated by the scaling factor 2^j, the pseudo frequency corresponding to the scale is expressed as:

$$F_s = \frac{F_c}{2^j \times \Delta},$$

where Δ is the sampling interval.

Taking the wavelet ‘coif5’ as an example, the centre frequency as seen from figure 5 is 0.68966 and thus the pseudo frequency corresponding to the scale 2^5 is 0.02155. As the sampling period is one month, the period corresponding to the pseudo frequency is 3.87 years.

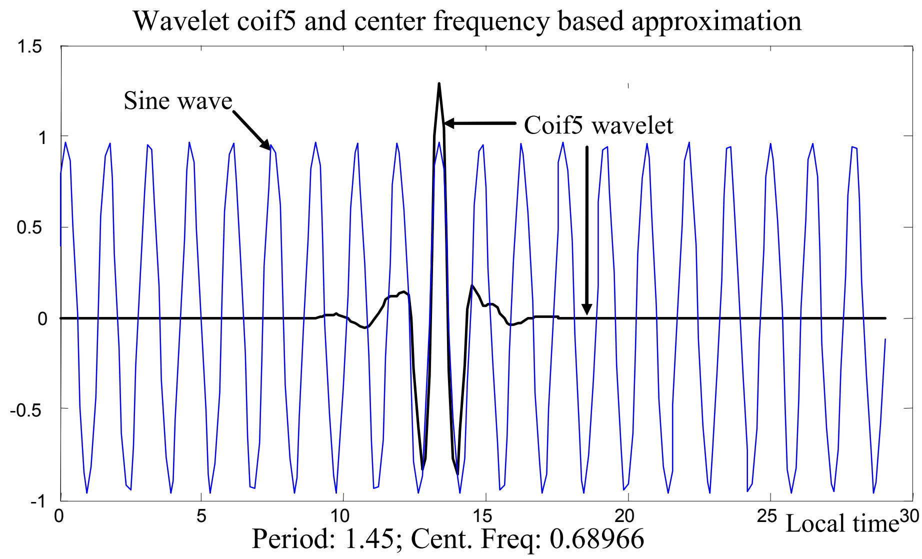

Figure 5: Scale in terms of equivalent sinusoidal frequency

Appendix B

Energy in returns versus levels

Table 3 is an energy table for returns as distinct from the levels (values) used in the text. As the table indicates, wavelet energy now concentrates in the high detail band, and there is little indication of any interior maximum or other sign of power at higher details.

Table 3: Energy decomposition for monthly asset returns

	NZ monthly return wavelet transform	US monthly return wavelet transform	JP monthly return wavelet transform	AU monthly return wavelet transform	USD/NZD forward monthly return wavelet transform
E(A7)(value)	0.0019	0.0029	0.0033	0.0030	0.0042
E(D1)	0.3843	0.1975	0.3582	0.1888	0.0946
E(D2)	0.1526	0.0778	0.1856	0.0827	0.0442
E(D3)	0.0887	0.0391	0.0916	0.0360	0.0268
E(D4)	0.0290	0.0168	0.0585	0.0223	0.0090
E(D5)	0.0110	0.0070	0.0325	0.0058	0.0055
E(D6)	0.0029	0.0124	0.0095	0.0025	0.0130
E(D7)	0.0069	0.0049	0.0036	0.0011	0.0026

Appendix C

Energy table for a geometric Ito process

Table 4 depicts the wavelet decomposition of a standard geometric Ito process with an annual drift of 10% and annual volatility of 20%. The simulated sample is a year with 365 daily observations for which the maximum scale for wavelet decomposition is 8. As is evident from the table, the energy increases along the scale level.

Table 4: Energy decomposition for geometric Ito process

Detail level	Detail Energy (%)
1	0.90%
2	1.12%
3	1.65%
4	7.29%
5	2.99%
6	8.61%
7	19.26%
8	58.18%

Appendix D

Data definitions and sources

Table 5: Data definitions and sources

Data	Definition	Resource
MSCI USA	USA stock total return index	MSCI
MSCI NEW ZEALAND	New Zealand stock total return index	MSCI
MSCI AUSTRALIA	Australia stock total return index	MSCI
MSCI JAPAN	Japan stock total return index	MSCI
NZD TO 1 USD	Spot exchange rate (USD as the commodity currency and NZD as the terms currency, 1USD=SNZD)	BBI
JPY TO 1 USD	Spot exchange rate (USD as the commodity currency and JPY as the term currency, 1USD=SJPY)	BBI
AUSTRALIAN $ TO US $	Spot exchange rate (USD as the commodity currency and AUD as the term currency, 1USD=SAUD)	BBI
NEW ZEALAND $ TO US $ 1MFWD	One month forward exchange rate (same expression as the spot rate)	BBI
NZ Bank Bill Rate	NZ 30 day bank bill rate	NZ Reserve Bank
US CD Rate	One month US certificate of deposit rate	Federal Reserve Bank
US Treasury Bond	US 10 year treasury bond	Federal Reserve Bank

References

Baillie, R.T., Long memory processes and fractional integration in econometrics. *Journal of Econometrics,* 1996, 73, 5–59.

Baillie, R.T., Bollerslev, T. and Mikkelsen, H. O., Fractionally integrated generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics,* 1996, 74, 3-30.

Barnett, W.A., Geweke, J. and Shell, K., *Economic Complexity: Chaos, Sunspots, Bubbles and Nonlinearity,* 1989, New York: Cambridge University Press.

Bhardwaj, G. and Swanson, N. R. An empirical investigation of the usefulness of ARFIMA models for predicting macroeconomic and financial time series. *Journal of Econometrics,* 2006, 131, 539-578.

Bowden, R. J., Spectral utility functions and the design of a stationary system. *Econometrica* 1977, 45, 1007-1012.

Bowden, R. J., *Statistical Games and Human Affairs*. 1988, New York: Cambridge University Press.

Bowden, R. J., The zero capital approach to portfolio enhancement and overlay management. *Quantitative Finance*, 2003, 3, 251-261.

Bowden, R. J. and Zhu, J., Asymmetric hedging of the corporate terms of trade. *Journal of Futures Markets*, 2006, 26, 1059-1088.

Bowden, R.J., The generalized value at risk admissible set: constraint consistency and portfolio outcomes. *Quantitative Finance,* 2006, 6, 159-171.

Bowden, R. J. and Zhu, J., The agribusiness cycles and its wavelets. *Empirical Economics*, 2007a, forthcoming, online ISSN 1435-8921.

Bowden, R. J. and Zhu, J., Which are the world's wobblier currencies? Reference exchange rates and their variation. *Studies in Nonlinear Dynamics and Econometrics,* 2007b, forthcoming.

Campbell, J. Y. and Viceira, L. M., *Strategic Asset Allocation: Portfolio Choice for Long-term Investors*. 2002, Oxford University Press, New York.

Candelon, B. and Gil-Alana, L. A., Fractional integration and business cycle features. *Empirical Economics*, 2004, 29, 343-359.

Capobianco, E., Multiscale analysis of stock index return volatility. *Computational Economics*, 2004, 23, 219-237.

Cohen, A., Daubechies, I. and Feauveau, J. C., Biorthogonal basis of compactly supported wavelets. *Communications in Pure and Applied Mathematics*, 1992, 45, 485-560.

Coifman, R., Meyer, Y., Quake, S., and Wickerhauser, V., Signal processing and compression with wavelet packets. Numerical Algorithms Research Group, 1990,Yale University, New Haven, CT, USA.

Crowley, P., An intuitive guide to wavelets. Bank of Finland/College of Business Texas A&M University, 2005, Patrick.Crowley@bof.fi

Dantzig, G. B., *Linear Programming and extensions*. 1963, Princeton University Press.

Daubechies, I., Orthonormal bases of compactly supported wavelets. *Communications on Pure and Applied Mathematics,* 1988, 41: 909-96.

Daubechies, I., The wavelet transform, time-frequency localization and signal analysis. *IEEE Transactions on Information Theory*, 1990, 36 (5), 961-1005.

Daubechies, I., Ten lectures on wavelets. *Society for Industrial and Applied Mathematics*, 1992, Capital City Press, Montpelier, VT, USA.

Gencay,R., Selcuk, F., and Whitcher, B., *An Introduction to Wavelets and Other Filtering Methods in Finance and Economics*. 2002, Academic Press, New York.

Gencay, R. and Selcuk, F., Asymmetry of information flow between volatilities across time scales, Econometric Society, 2004 North American Winter Meeting.

Granger, C. W. J., Long memory relationships and the aggregation of dynamic models. *Journal of Econometrics*, 1980, 14, 227–238.

Guidolin, M. and Timmermann, A., Strategic asset allocation and consumption decisions under multivariate regime switching. Proceeding Paper No. 349, 2004, Econometric Society 2004 Australasian Meetings.

Hall, V.B. and McDermott, C.J., The NZ business cycle: return to golden days? Research Paper, School of Economics and Finance, 2006, Victoria University of Wellington.

Haugh, M. and Lo, A., Asset allocation and derivatives. *Quantitative Finance*, 2001, 1, 45-72.

Hong, Harrison, Scheinkman, J. and Xiong, W., Asset float and speculative bubbles. *Journal of Finance*, 2006, LXI, 1073-1117.

Ibbotson, R.G. and Sinquefield, R., Stocks, bonds, bills and inflation: Year by year historical returns (1926-1976). *Journal of Business*, 1976, 49, 11-47.

In, F. and Kim, S., The hedge ratio and the empirical relationship between the stock and futures markets: a new approach using wavelet analysis, *Journal of Business*, 2006, 79,799-820.

Kim, K., Buckle, R.A. and Hall, V.B., Dating New Zealand business cycles. *New Zealand Economic Papers*, 1995, 29,143-171.

Kohn, R.V. and Papazoglu-Statescu, O.M., On the equivalence of the static and dynamic asset allocation problems. *Quantitative Finance*, 2006, 6, 173-183.

Kojima, T., Speculative bubbles and crashes in stock markets: an interacting-agent model of speculative activity. *Physica A*, 2000, 287, 493-506

Lance, F., Sources of exchange rate and price level fluctuations in two commodity exporting countries: Australia and New Zealand. *Economic Record*, 1996, 72, 345-358.

Lee, H., International transmission of stock market movements: A wavelet analysis. *Applied Economics Letters*, 2004, 11, 197-201.

Lux, T., Herd behavior, bubbles and crashes. *Economic Journal*, 1995, 105, 881-896.

Mallat, S., A theory for multiresolution signal decomposition: the wavelet representation. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 1989, 11, 674-93.

Misiti, Michel, Misiti, Y., Oppenheim, G. and Poggi, J. M., *Wavelet Toolbox user's guide*. 2005, www.mathworks.com.

Murty, K. G., *Linear and Combinatorial Programming*, 1976, John Wiley.

Otrok, C., Spectral welfare cost functions. *International Economic Review*, 2001, 42, 345-367.

Percival, D. P., On estimation of the wavelet variance. *Biometrika*, 1995, 82, 619-623.

Percival, D.B., and Walden, A.T., *Wavelet Methods for Time Series Analysis*. 2000, Cambridge: University Press.

Ramsey, J. B., The contribution of wavelets to the analysis of economic and financial data. *Phil. Trans. R. Soc. Lond. A* , 1999, 357, 2593-2606.

Ramsey, J. B., Wavelets in economics and finance: Past and future. *Studies in Nonlinear Dynamics & Econometrics*, 2002, 6, No. 3.

Rockafellar, R.T., Duality in nonlinear programming. In *Mathematics of the Decision Sciences Part I*. 1968, (ed) G.B. Dantzig and A.F. Veinott Jr., Providence R.I.:American Mathematical Society.

Ross, S. A., The arbitrage theory of capital pricing. *Journal of Economic Theory*, 1976, 13, 341-360.

Scheinkman, J. A. and Xiong, W., Overconfidence and speculative bubbles. *Journal of Political Economy*, 2003, 111, 1183-1219.

Schleicher, C., An introduction to wavelets for economists. Bank of Canada Working Paper 2002-3.

Solnik, B., An equilibrium model of the international capital market. *Journal of Economic Theory*, 1974, 8, 500-524.

Thorp, S., That courage is not inconsistent with caution: Currency hedging for superannuation funds. *Economic Record*, 2005, 81, 38-50.

Vuorenmaa, T. A., A wavelet analysis of scaling laws and long-memory in stock market volatility. Discussion Paper 27, 2005, Bank of Finland Research.