Asymmetric hedging of the corporate terms of trade

Roger Bowden and Jennifer Zhu*

Abstract

Risk management techniques such as value at risk and conditional value at risk focus attention on protecting the downside exposures without penalising the upside. The implied welfare functions are equivalent to an otherwise risk neutral agent with a put option exposure on the downside. The correspondence can be exploited to design smoother loss measures and numerically based solutions for optimal hedge ratios. A statistically well-adapted hedge object for the firm is the corporate terms of trade, which balances up output and expense prices as a single index related to the net profit margin. The methods are applied to the NZ dairy industry to derive optimal foreign exchange forwards based hedges. It is not always optimal to rely solely on forward discounts or premiums.

Key words: Conditional value at risk; forward foreign exchange; hedging; option equivalent utility functions; terms of trade.

JEL numbers: G10,G11,G15; F31,F37; M21;Q13

* The authors are both with the School of Economics and Finance, Victoria University of Wellington, PO Box 600 Wellington, New Zealand; emails roger.Bowden@vuw.ac.nz;jennifer.zhu@vuw.ac.nz.
Thanks go to Jin Seo Cho, Dawn Lorimer and an anonymous JFM referee for comments that helped to focus discussion. The paper was delivered by Jennifer Zhu at the Australasian economics and finance PhD students conference in November 2005, and both authors are especially grateful to discussant Susan Thorp of UTS, for a picking up some slips and other helpful comments.

I Introduction

The nature of the objective function in problems of hedging and portfolio management has become highlighted in recent attempts to encompass asymmetric loss via value at risk and conditional value at risk. The former (e.g. JP Morgan 2001, Jorion 2000) essentially refers to lump sum type penalties if the outcome return or value falls below an acceptable region; while conditional value at risk or CVaR refers to the disutility to be attached to the left hand tail as a whole (Urysaev 2000, Palmquist *et al*. 2001, Rockafellar *et al*. 2002, Alexander and Baptista 2004). On the other hand, an exclusive concern with the hazard region ignores the benefits to be gained from good outcomes, which in many cases amounts to core business for the firm or fund. A more embracing formulation would allow for both, namely the expected payoff in the good zone and the expected damage in the bad zone, weighting the two according to managerial preferences. This amounts to specifying the underlying utility function to have segmented or smoothed segmented form, with a recognisable nesting within it of value at risk or conditional value at risk. Desired value at risk constitutes the critical point, which becomes a natural calibration point for the resulting utility function.

The asymmetry problem becomes acute in hedging. In this context, the manager definitely wants higher returns or cash flows in the good zone, and might very well be taken as effectively risk neutral in this zone. But the point of hedging, from a corporate finance perspective, is to minimise the extreme value hazard on the downside. The hazard can lead to bankruptcy or other costs, the risk of which diminishes the value of the firm, as well as creating corporate and personal stress. Or the downside can involve depositor outflow, a risk for any funds manager. The problem is how to establish hedging decision rules that balance up the upside and downside. The present paper shows how this can be done in a manner consistent with traditional Von Neumann -Morgenstern decision theory under risk, giving the manager some guidance as to how to calibrate organisational risk preferences into the objective function. The most basic version entails a piecewise linear utility function with penalty slope beyond a chosen VaR point. The penalty slope – which generates the conditional value at risk – can be calibrated by presenting the manager with the conditional outcomes in the two zones. The effect is as though an otherwise risk neutral manager has written a put option with strike price set at the specified value at risk critical point.

Operational versions that cater for convergence difficulties associated with a finite sample size can be devised by ‘fuzzifying’ the switch point, drawing in this respect on the ideas of Zadeh (1965, 1970). However, we suggest a new smoothing procedure that takes advantage of the options interpretation of the utility function. The corresponding Black-Scholes price provides the smoothing, and the outcome can be taken as a utility function in its own right, in manner similar to the bankruptcy cost justification for hedging in corporate finance, which amounts to being short an option on the costs of financial distress. The resulting hazard-modified expectational approach yields a flexible functional form catering for different risk reward preferences, while preserving value at risk type critical points as calibration references. One can also use ordered mean difference methodology to explore the robustness of the resulting hedges to utility functions that may be of yet more general form.

The chosen context, namely foreign exchange-based hedging, highlights the nature of the decision problem as between expected reward and risk. Some countries have interest rates persistently higher than their trading partners, e.g. capital importing countries like Australia or New Zealand. This implies a forward discount on their exchange rates and a natural motivation for goods and services exporters to sell forward their foreign exchange receipts. A popular view has grown up in such countries that exporters should always and fully utilise the forward rate to convert their foreign currency earnings. Capital exporting countries like Japan or Switzerland have a complementary motivation for goods and services importers. However, exclusive use of forwards amounts to paying attention only to the expected positive return on the forward conversion rate relative to the spot rate at the time, ignoring the possibility of the eventual spot rate conversion error being very large in the wrong direction. A more appropriate portfolio might contain both the forward and the unhedged spot, or a range of forward maturities, in proportions that depend upon managerial preferences as between gains if things go right and losses if things go bad.

Addressing problems of this kind also requires attention to the hedging data framework. From the corporate finance point of view, a complete solution would entail hedging total free cash flows to investors. Operationally, this is a formidable problem for a number of reasons, including data availability frequency, and the issue of normalising the variables to express them on a common basis, e.g. avoiding the problem of hedging stationary with non stationary time series data. We suggest a key indicator approach that isolates the most sensitive exposures in the form of the corporate terms of trade, as a proxy for corporate profitability in the form of the net profit margin. The chosen application is to the hedging of dairy farmer incomes and cash flows, a major exposure for the New Zealand economy and one that has produced adverse hedging outcomes of major proportions in the recent past. Most of the exposures and survivability threats can be summarised in a single industrial terms of trade index which expressed farmer income relative to farmer costs, somewhat analogous to the real exchange rate for the economy at large. The major factors (in VaR terminology) are the exchange rate, world commodity prices and farmer expense prices. The chosen market-based hedge is the exchange rate, in the form of forward rate transactions. As a hedge instrument this is incomplete, for there will be basis risk from the independent movement of input prices or commodity prices, although both the latter will also show some correlation with the exchange rate.

The scheme of the paper is as follows. Section II develops the basic theory of the risk-averse objective function, developed as a simple analogy to the conditional value at risk criterion. The point of departure is taken as simple expected value maximisation, i.e. risk neutrality, over most of the range. Risk aversion arises because of sensitivity to lower tail outcomes, so that in effect the manager or investor has written a put option in favour of nature in this zone. Calibration issues are discussed, and the smoothed version established in terms of the price of the implied put option. This is followed by a short discussion of the nature of hedging and the accounting framework. Section III describes the chosen context, establishing the farmer terms of trade and its relation with the net profit margin as the hedge object. The objective function is calibrated off historical data, which also illustrates the advantages of forwards and the potential risk in using them. The section concludes by setting up the decision problem in a form amenable to numerical computation. Section IV describes the results, in terms of both shorter and longer run hedging. Section V contains some general observations and recommendations.

II Hedging decision theory

2.1 The objective function

Let R stand for an exposure outcome: a portfolio return, value, cash flow, or some other dimension of corporate or managerial exposure. In the application of the present paper it is taken as the net profit margin, but the context of the present section can be quite general. In order to highlight the asymmetry aspect, it is assumed that over most much of the range of R, the manager is risk neutral, but over lower ranges risk neutrality is replaced by risk aversion. A motivation for this particular formulation arises from the value at risk literature (see below) though it should be stressed at the outset that not all agents will have this form of welfare function. For instance, empirical studies of investor behaviour commonly suggest coefficients of relative risk aversion of around 2, which are perhaps more consistent with decreasing marginal utility over upper as well as lower ranges. On the other hand, the importance of bonus-seeking behaviour in managerial life on the upside, coupled with the prospect of being fired on the downside, might produce a risk aversion profile more of the type canvassed in the present contribution, which can also admit further powering up of downside aversion for very bad outcomes. In what follows, we shall initially motivate discussion with reference to the extended value at risk literature, but the options interpretation that results survives some of the difficulties associated with the value at risk framework and allows beneficial upside exposure. However, a virtue of value at risk is that it remains a useful guide to calibration of the resulting utility function, so it is a good place to start.

The value at risk (VaR) refers to the critical point for the lower tail (e.g. 5% or 10%) of the distribution of R . Originally in the nature of a portfolio 'diagnostic', and still much used in that way in the regulatory context, it later became the basis for a portfolio choice theory, which in turn can be formulated in several different ways. One version seeks to minimise the value at risk for a given significance level (critical probability), and subject also to given minimal performance on other metrics such as the expected value. Other versions set at the outset a given value P to be the VaR critical point and seek to maximise the expected portfolio value outcome, subject to the requirement that the probability of falling short of P is less than or equal to the chosen significance value (5% or 10%, etc.). In effect, VaR provides the risk aversion. The latter version amounts to asserting that the firm's core business is to maximise expected value, but that the VaR critical point P now provides the tolerance level for adverse outcomes. The virtue of a pre-set critical point P, as opposed to a pre-set critical probability, has been underscored by Alexander and Baptista (2004). They point out that for any given critical probability, a low standard deviation portfolio will in some cases have a higher value of P than a high standard deviation portfolio, giving the former an increased prospect of adverse events in the critical region less than P. Designating the critical point P as a calibration parameter will be the general approach followed in the present paper.

In addition, conditional value at risk (CVaR) imposes a further requirement that the expected shortfall $P - E[R|R \le P]$ is to be less than some given tolerance number ν, say; (Urysaev 2000, Palmquist et al 2001). The CVaR approach is concerned with the expected tail length in the form of the truncated or censored expectation, $E[R|R \le P]$; one is requiring that this is not too low. The motivation comes from the long left hand tail possibility, one that may do considerable damage to managerial welfare in a manner that would not be captured by the VaR limit just by itself (Basak and Shapiro 2001). A difficulty is that imposing both VaR and CVaR as formal constraints results in an implied utility function that is inconsistent with Von

Neumann-Morgenstern utility theory – the effective utility function is not uniformly increasing around the critical point and may indeed spike upwards at $R=P$ (Bowden 2006; see also Artzner *et al* 1997,1999 for operational consequences).

However, a convenient way to capture the essence of the CVaR approach is to specify at the outset a utility function of the form

$$U(R;P) = R - P + \beta \min(R - P, 0)\,.$$

For later convenience this can be written as

$$U(R;P) = R - P + \beta(R - P)SF(P - R)\,, \tag{1}$$

with the unit step function defined by $SF(x) = \begin{cases} 1 & \Leftrightarrow x > 0 \\ 0.5 & \Leftrightarrow x = 0 \\ 0 & \Leftrightarrow x < 0 \end{cases}$.

In the formal CVaR programming problem, the objective function is locally equivalent to (1) with the constant β emerging as the shadow price of the constraint, evaluated at the optimal solution (Bowden 2004). However the constraint value has itself to be allocated by the user as the constant ν, so that an alternative is to have the user set β at the outset. The constant P is set as the desired value at risk critical point.

Figure 1 illustrates the objective function. Above the critical point P, utility increases linearly with the return or value outcome. Below the critical point P, the slope increases from unity to $1+\beta$, indicating a higher disutility as the return sinks below the critical point; this disutility can also be 'powered up' if necessary (see below).

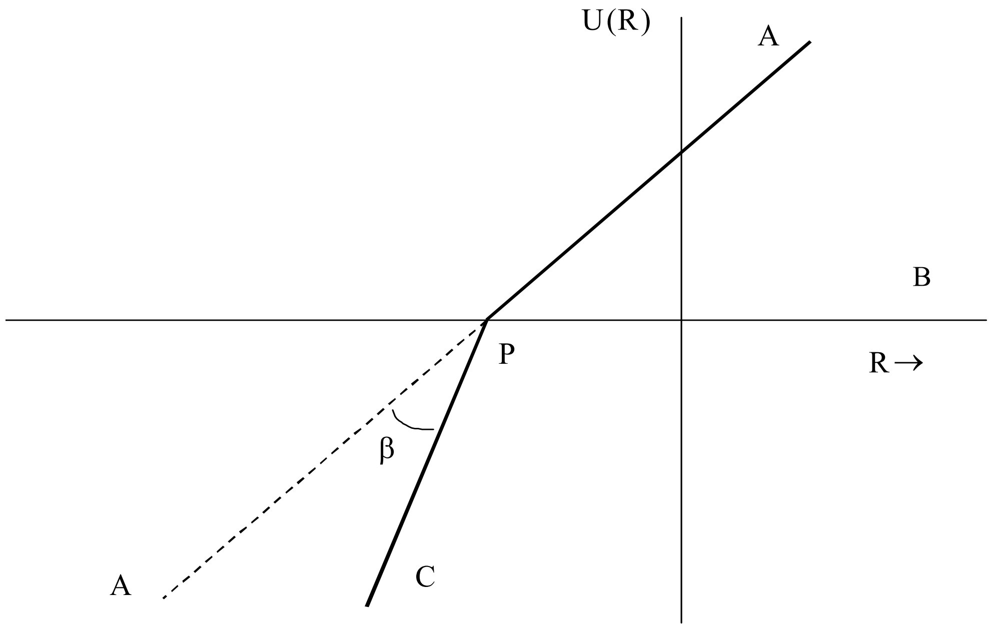

Figure 1: Objective function

The effect of figure 1 is as though an otherwise risk-neutral manager – with a linear utility function - has been compelled to write β put options with strike price at P, with payoff profile given by APC. The contingent cost of doing so increases as P increases and also as the number β increases. The natural position effectively amounts to a portfolio of long a forward contract with price P and short β put options with P as the strike price. One could also encompass VaR-constrained programming within the option equivalent framework, by noting that there are implied lump sum penalties for exceeding VaR limits, either explicitly so, or else emerging from the duality properties of the optimisation subject to a given VaR limit (Bowden 2006). This amounts to being short in a binary option. The options approach corresponds to a common rationale for hedging from the theory of corporate finance, namely the avoidance of bankruptcy costs, which is as though the managers have written put options on the firm's value in favour of third party claimants (see section V).

The connection with CVaR may be viewed by taking the expected utility:

$$E[U(R;P)] = E[R] - P + \beta F_R(P) \times E[R - P | R \leq P], \tag{2}$$

where $F_R(P)$ is the distribution function of R evaluated at $R = P$. If one chooses P to satisfy a pre-set VaR critical point then $F_R(P)$ becomes a constant, and the risk penalty is provided by the conditional value at risk term $E[R - P | R \leq P]$. The parameter β is set as a penalty to the CVaR component relative to the overall or unconditional expectation $E[R]$; its value will incorporate or otherwise be influenced by the desired VaR critical point (see below).

A further possibility is to 'power up' the bad zone disutility as:

$$U(R;P) = R - P - \beta \, | R - P |^{\kappa} \; SF(P - R) \tag{3}$$

where the constant $\kappa > 1$. This can cover cases such as those pointed out by Yamai and Yoshiba (2002), where the expected shortfall fails to eliminate the tail risk beyond a specific threshold (i.e. very long tail effects). Of course, it could also be the case that marginal disutility is increasing with larger losses. In term of option equivalence, the effect would be as though the manager had written power options on the downside. Note that the special case $\kappa = 2$ and $P = E[R]$ would correspond to using the semivariance as a downside risk metric. Ferreira and Goncalves (2004) also employ a powering up metric, though they use the conditional κth order tail moment; the conditionality is similar to omitting the factor $F_R(P)$ in expression (2).

Calibration

Risk aversion depends upon the two parameters P and β, so the user has to calibrate their values at the outset.

(a) The discomfort point P could be calibrated by examining the firm's cash flows, in the same way that ones does with cash flow at risk. In terms of the present application, a given terms of trade or net profit margin implies a cash flow for the operation, and danger points for the latter generate corresponding danger points for the former. Alternatively, one can plot the natural exposure (unhedged terms of trade) over the past and isolate the lower 5 or 10% crucial points, identifying these as alternative values of P. This would amount to saying that the past natural or unhedged values constituted a desired range for the hedged as well as unhedged exposures. If that is not true, one could adjust P upwards, e.g. identifying the desired hedged 5% VaR with the natural or unhedged 10% VaR. Section 3.2 below employs a combination of these approaches.

(b) The risk aversion parameter β could be calibrated by rearranging expression (2) as:

$$E[U(R;P)] = \omega_L \, E[R-P|R \le P] + \omega_U \, E[R-P \mid R > P]. \tag{4}$$

where $\omega_L = (1+\beta)F_R(P)$ and $w_U = 1-F_R(P)$.

Expression (4) splits up the expected utility into the conditional expectation contributions respectively less and greater than the critical point P (the 'bad' and 'good' zones). The bad zone expectation (i.e. the first on the right hand side) is effectively the conditional value at risk. Simple risk neutrality would correspond to $\beta = 0$. On the other hand, the manager might wish to weight the two zone expectations equally, setting off the expected gain on the upside against the conditional value at risk. Setting $\omega_L = \omega_U$ gives $\beta = (1-2F_R(P))/F_R(P)$, so if P is chosen as the 10% point, then $\beta = 8$; or if P is the 5% point, $\beta = 18$. Increasing β beyond these values amounts overweighting the conditional value at risk relative to the expected gain on the upside, so this is greater risk aversion. The extreme case $\beta \to \infty$ is interpreted to mean that the manager is concerned only with the conditional value at risk $E[R-P|R \le P]$.

Other methods of calibration are available. For instance, one could compute a risk aversion coefficient such as the generalised Rubinstein risk premium θ (e.g. Bowden 2005) which says that expected utility is certainty-equivalent to an outcome of $\mu_R - \theta$. As θ depends on the aversion parameter β, one can set the value of beta to accord with any desired certainty-equivalent outcome. Appendix A gives the relevant formula for the GR premium in the present context.

Whatever the method used, a reality check is always useful where calibration is concerned. We suggest running the optimisation with a range of beta values and computing risk diagnostics for each. For instance, there is an inverse relationship between beta and the conditional value at risk, so if the latter is seen as too high, then the remedy would be to increase the value of beta. Likewise, it would seem a useful precaution to explore the effects of different powering constants (κ) as in expression (3).

Smoothing the objective function

In practice, it is useful to replace the kink at $R=P$ in figure 1 with a smoother utility transition while preserving the point P as a useful calibration benchmark. This could be seen as preserving a closer relationship with traditional expected utility theory, which would view risk aversion behaviour as psychological in nature, rather than determined by administrative or regulatory necessity, which can result in discontinuity. However there are reasons to smooth the transition even where special weight has to be placed is on the VaR point P, as in the administrative approach. A first reason is computational: one often experiences convergence difficulties with the optimisation of expected step or segmented (ramp style) functions, where the sample expectation has to be estimated off a finite number of observations. In the present context, there will be fewer observations in the critical zone ($R \le P$), and small changes in the trial solution for the proposed hedge ratios may have no effect, leading to the numerical optimisation programme getting stuck. The problem is well known in switching regression theory, where there may be very few observations at or around the chosen switch point. Secondly, it could be remarked that if the payoff function resembles that of a put option, as earlier pointed out, then one could consider the disutility as in the nature of a cost related to the price of the option, which is inherently continuous in nature.

Alternative ways of smoothing can be derived, one based on fuzzy logic and the other drawing on the options interpretation. The former operates by replacing the implicit step function in (1) with a fuzzy membership function (Zadeh 1965, 1970) similar to a notional probability distribution, adjusting the variance parameter to achieve different degrees of smoothing. Appendix B describes this method in more detail. It is very easy to implement and is correctly centred at the critical point P, but has the disadvantage that it can assume a local negative slope for very high values of the aversion parameter β.

The options-based method is more appealing in terms of the underlying economic interpretation. Given any value R, we price the implied put options (β of them) by assuming the current physical value is R and the strike price is P. Thus if the outcome is R, one can think of the manager as being assigned a penalty equal to the value of a put option at that price, with strike set at P. The market price of the option is treated as the risk aversion penalty. One could assign the volatility parameter as that of the unhedged position. In terms of corporate finance theory, the objective would be to limit the value of the options held by third party claimants in the event of bankruptcy.

Given an assumed volatility number σ, the Black Scholes put option price is inserted into expression (1) in place of the ramp function element $\beta \min(R-P,0)$, to give:

$$
\begin{aligned}
&U(R;P) = R - P - \beta \times \pi(R;P), \\
&where: \quad \pi = e^{-r\times T} \times [P \times N(-d_2) - R \times N(-d_1)] \qquad (5) \\
&d_1 = \frac{\ln(\frac{R}{P}) + (r + \frac{\sigma^2}{2} \times T)}{\sigma\sqrt{T}}, d_2 = d_1 - \sigma \times \sqrt{T} \quad .
\end{aligned}
$$

Setting the parameters σ, *r* and *T* depends on interpretational aspects. As earlier suggested, one could assign σ as the volatility of the unhedged position, *r* as the market rate and *T* as the length of a performance evaluation horizon. On the other hand, if all that is required is a smoothing device, then one sets σ as small as is consistent with computational convergence requirements, and might as well fix *r = 0* and *T = 1*. Figure 2 shows what happens, in this case for the application of section III, choosing σ *=0.025, r =0 and T =1*.

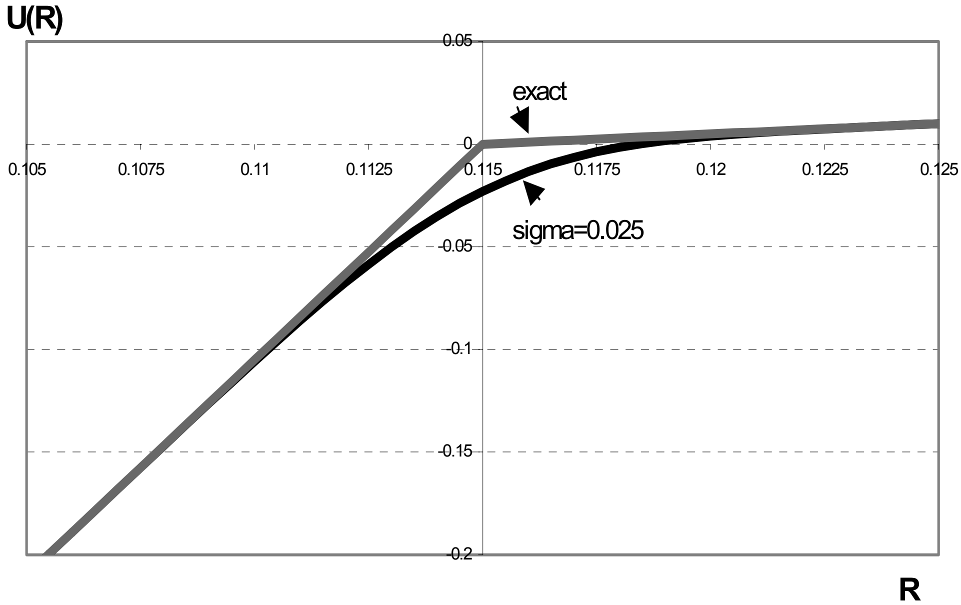

Figure 2: Option equivalent approach to the exact utility function

[Note: To preserve the positivity needed for the logs in the B-S price, a constant has been added to the natural terms of trade variable, though P is still formally located as the zero point].

2.2 Rolling hedge frameworks

Academic discussions are commonly confined to a simple forward- spot hedge, using a single forward contract of pre-set maturity. On the other hand, many practical hedge rules are of the rolling hedge variety, staggering the cover over a variety of forward maturities in order to achieve a smoothing of the exposure. This raises some measurement problems that must be resolved if one is to adequately compare hedge outcomes or design optimal hedge rules.

To illustrate, suppose that a series of forward rate contracts, at prices collectively denoted *F*, are to be used to hedge a foreign exchange rate exposure, denoted by *S*. (The precise relationship of this to the objective outcome *R* of the previous section will be discussed in 2.3 below). Consider, for instance, a common rolling quarterly hedge of the following form in foreign exchange (FX) risk management: Hedge 100% of Q1 exposure, 75% of Q2 exposure, 50% of Q3 exposure and 25% of Q4 exposure. In other words, one makes sure that the near quarters are hedged more completely than the further out, constantly adjusting the proportions upwards as the distant quarter approaches. Once a steady state has been achieved, the hedge portfolio for each real time period will consist of equal proportions of one quarter, two quarter, three quarter and four quarter forwards. Table 1 shows how things will look. The firm has a constant physical exposure of $100 foreign dollars (or euros etc.) each quarter.

Real or conversion time ↓	Time to maturity →	t+1	t+2	t+3	t+4	t+5	t+6	t+7	......
......									
t		$25.00	$25.00	$25.00	$25.00				
t+1			$25.00	$25.00	$25.00	$25.00			
t+2				$25.00	$25.00	$25.00	$25.00		
t+3					$25.00	$25.00	$25.00	$25.00	
[t+4]					[$100]				

Table 1: Alternative hedge measurement frameworks

Under the *effective conversion rate* (ECR) approach we concentrate on the vertical forward rates for any given real time. Thus at time *t+4*, the ECR achieved will be a function of the forward rates at times *t, t+1, t+2,* and *t+3* as in the grey shaded column. The ECR can be compared directly with the actual spot rate at the time S_{t+4}.

Under the *mark to market approach*, it is the horizontal rows that are of interest. Thus at any real time *t*, one is holding an portfolio of $25 worth of one, two, three and four period ahead contracts. It is the change in values of these contracts between times *t* and *t+1* that constitutes the cash flow, when taken together with the physical transaction at time *t+1*. This is the way that mark to market profits or losses are recorded in most treasury environments.

Appendix C contains a more detailed comparison of the two accounting methods and how they relate to the cash flow from rolling the hedges forward. The ECR method has the advantage that the conversion rate achieved can be compared directly with the spot rate for the physical transaction. Thus in the above example, we have

$$ECR_{t+4} = 0.25F_{4,t} + 0.25F_{3,t+1} + 0.25F_{2,t+2} + 0.25F_{1,t+3}$$

where $F_{m,n}$ = forward price for maturity m, bought at time n. The effective conversion rate ECR_{t+4} can then be compared directly with S_{t+4}, the spot rate, which provides a way of evaluating and comparing hedge strategies.

More generally, one could have arbitrary four quarter rolling hedges of the form

$$ECR_{t+4} = h_1F_{4,t} + h_2F_{3,t+1} + h_3F_{2,t+2} + h_4F_{1,t+3} + (1-h_1-h_2-h_3-h_4)\times S_{t+4}, \qquad (6)$$

where $h_1 \geq 0$, $h_2 \geq 0$, $h_3 \geq 0$, $h_4 \geq 0$ *and* $h_1+h_2+h_3+h_4 \leq 1$.

This would amount to leaving a proportion of the exposure unhedged, i.e. exposed to the final spot exchange rate S_{t+4}. In effect one is constructing a portfolio of forward rate contracts combined with some unprotected spot exposure.

2.3 The decision problem

The final task is to bring the two strands together, namely the objective function and the hedging strategy. It is helpful to adapt some terminology from the optimal control literature, of which hedging can be regarded as an instance. By a *hedge instrument* we shall mean a market-based contract that may be bought or sold in the pursuit of risk management. By a *hedge target* we will mean a variable that the hedge instrument is naturally adapted to. By a *hedge object,* we shall mean a variable that enters the hedging objective function as the ultimate object of control; it will typically depend in some way on the hedge target variable.

Thus the hedge object is the variable *R* that appears in the objective function of section 2.1. In turn, *R* might depend wholly or partly on a foreign exchange rate *S*. The latter is a suitable hedge target because of the availability of a set of FX forward contracts adapted to it; the latter are the hedge instruments. In the application that follows, the objective is to hedge the company's net profit margin as represented by the log terms of trade, which depends among other things on the exchange rate. The object *R* is the log terms of trade and the hedge target will be the exchange rate. The hedge instruments are exchange rate forwards and by using these one replaces the spot exchange rate as it appears in the industrial terms of trade with the effective conversion rate. By doing so, the distribution of the object variable R comes to depend upon the set of hedge weights h_i .The objective is then to choose the optimal hedge weights h_i in order to maximise the expected value of the objective function.

III Application

The dairy industry is one of the underpinnings of New Zealand economic activity, and the second largest earner of foreign exchange after tourism. The dairy company Fonterra is NZ's largest company, and like the other smaller dairy companies, is a cooperative owned by its farmer suppliers. Many of the suppliers are themselves significant private companies, milking thousands of cows. All NZ dairy farmers inherit the same problem, of risk management in the face of an uncertain environment, with volatile cash flows and incomes as a consequence.

A major source of uncertainty is the exchange rate. The bulk of NZ dairy produce is contracted in terms of US dollars, and the USD/NZD exchange rate is consequently the major single exposure. It has been highly volatile, trading over the last ten years in a range of about 1.35 to 2.55 USD/NZD. Dairy commodity prices are also a significant source of variation, correlated to some extent with the exchange rate. In addition, farmers are exposed to cost variations via the local economy and once again via the exchange rate for equipment, fertiliser, chemicals etc. In order to cope with this volatility, companies such as Fonterra have been highly active in the FX hedging market, not always with good results. Thus Fonterra's precursor the NZ Dairy Board took a charge of almost NZD500 million against its 1998 accounts following a hedge gone wrong (Bowden 2005, ch12). The design of effective hedging rules is therefore an important task for the industry.

3.1 The hedging objective

Hedge policy is better formulated and understood in terms of single line hedging, which in the present context means finding a single objective variable to index or capture the intrinsic exposure to environmental variations. The need to do this has arisen in other contexts for the industry such as livestock breeding objectives. One such measure that appears suitable for the purpose is the farmer terms of trade, which effectively constructs a home currency price of output or income relative to the price of inputs. Taking logs gives a metric that is approximately equal to the net profit margin for the enterprise.

A schematic decomposition is as follows:

$$\begin{aligned} profitability\ index &= \frac{output\ price \times exchange\ rate \times output\ quantity}{\exp enses\ price \times input\ quantity} \\ &= [\frac{output\ price \times exchange\ rate}{\exp enses\ price}] \times [\frac{output\ quantity}{input\ quantity}] \\ &\sim terms\ of\ trade \times productivity. \end{aligned}$$

Productivity is treated as exogenous in what follows, though this should not serve to minimise the longer-term imperative to develop productive capacity as a defensive response to adverse price variation.

The above profitability index (PI) can be regarded as the ratio of operating revenue to costs. A more standard accounting ratio is the net profit margin (NPM), which is the ratio of revenue net of costs to total revenue. The relationship between the two is $PI = 1/(1 - NPM)$, so $\log(PI) \approx NPM$. The net profit margin is a widely used measure of managerial performance, and a possible target for bonus fixing in the case of large scale farming operations, which rely on a professional farm manager. A further advantage of using the log terms of trade is that it allows a log linear separation of the variables. Operationally, therefore, the hedge object will be taken as:

$$R = \log[\frac{dairy\ product\ price\ index \times effective\ exchange\ rate\ USD / NZD}{NZ\ farmer\ \exp ense\ price\ index}] \tag{7}$$

The NZ farmer expense price is used as an index of the farmer input prices (see section 3.5 below on data). In the absence of an input price series before 1992, we estimate it from the CPI index based on the regression result of farmer expense price index against the CPI index from 1992 to 2004. The output price index series is constructed by using the dairy product price index together with the corresponding exchange rate. The use of the farmer expense price index could alternatively be interpreted as hedging real farmer incomes.

3.2 Unhedged behaviour

Figure 3a shows the history of the farmer log terms of trade over the entire sample period. The variability is self-evident. Two possible discomfort points are marked in, corresponding to the 5% and 10% lower quantiles (P_5 and P_{10}). These points are also marked in on figure 3a as horizontal lines at $P_5 = -1.8739$ and $P_{10} = -1.8501$. Figure 3b depicts the same data in the form of a formal histogram (this should not be taken too literally as the data are not serially independent and do not need to be in what follows). There are appreciable tails both to the left and the right, with a suggestion of a secondary right hand mode. The latter would make simple variance reduction an economically suboptimal hedging decision criterion, throwing out the good times baby with the bad times bathwater.

Early 1997 was the decision point for the NZ Dairy Board's ill-fated hedging decision (Bowden 2005 ch12), in the course of which farmer bankruptcies were cited by the Board as a possible outcome if the NZD exchange rate strengthened any further. This suggests that the discomfort or pain point for the farmer terms of trade would have been about the 10% point (P_{10} = -1.85 in the above diagrams). This will be the maintained assumption, though the 5% point was also be used to illustrate the effect of relaxation, and the 20% point for a tightening.

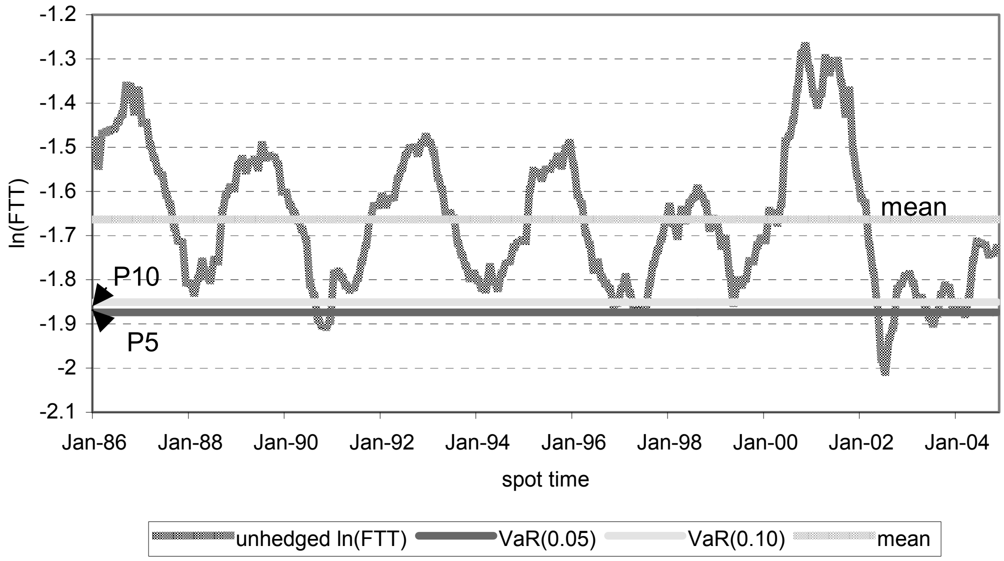

Figure 3a: unhedged farmer log terms of trade, time series

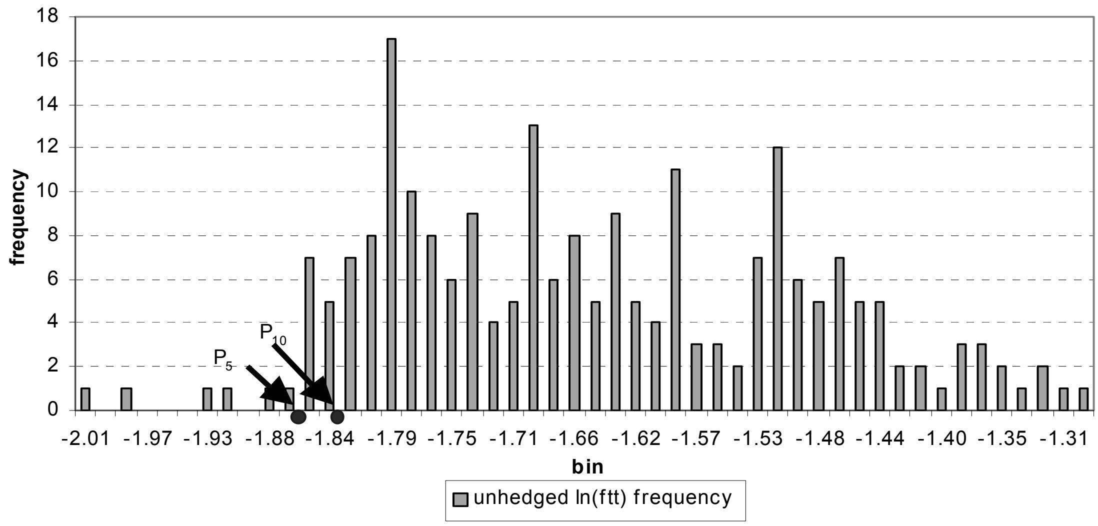

Figure 3b: Frequency histogram, natural exposures

3.3 Nature of the hedge

Direct hedges of dairy commodity prices – as farm-gate prices to the NZ industry – are not available. The two major sources of variability in the farmer terms of trade are the exchange rate and dairy commodity prices, with input prices more stable, at least in recent years. Areas of major discomfort are where low commodity prices coincide with a strong NZD/USD exchange rate. The effective hedge instrument for the farmer terms of trade is based on the forward exchange rate, which has both a direct relationship with the exchange rate and an indirect relationship with commodity and input prices, given that the NZD is widely regarded as a commodity currency.

As earlier mentioned, with a capital importing country like NZ there is a natural bias for exporters to make use of the forward exchange rate. NZ interest rates are chronically high relative to other OECD countries and this is reflected in the one year forward NZD discount plotted as the forward premium USD/NZD in figure 4.

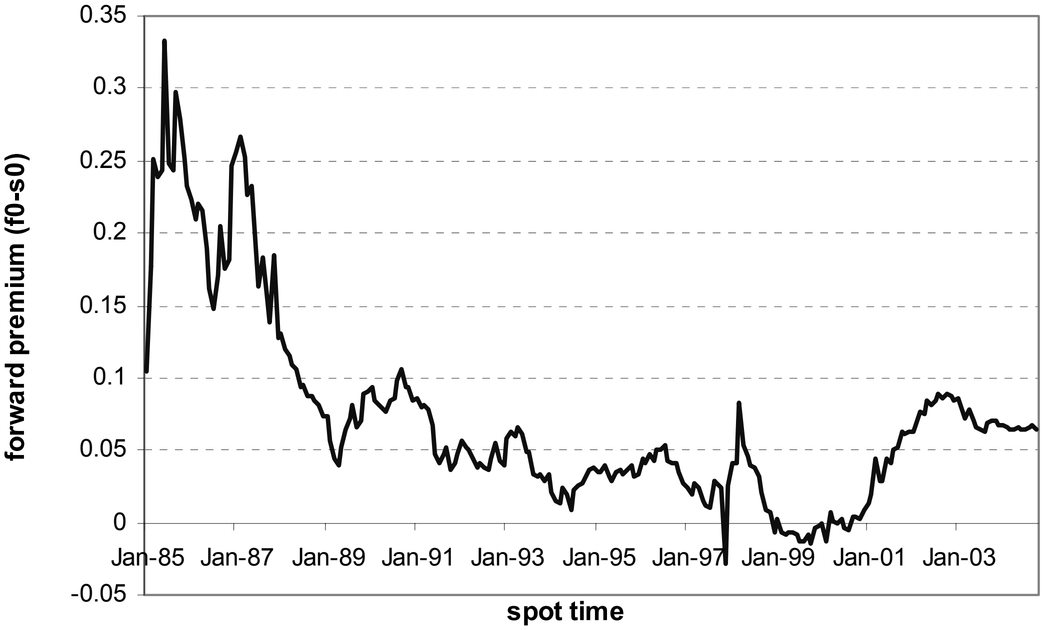

Figure 4: Forward premium on the USD/NZD

On the other hand, a forward discount on the home currency does not necessarily mean a forward profit, in terms of the actual conversion rate on maturity. Figure 5 graphs the difference between the forward rate at maturity and the spot rate as of the same day, presented in the form of a profit to the forward user. Use of the one-year forward would have resulted in a loss at times, as the Dairy Board discovered in their 1997 hedging programme. Likewise, use of the three-year forward (not depicted) is even riskier; consistent losses would have resulted between June '98 – June '02. Conversely, significant gains would be made where a fall in the USD reinforced the forward rate discount, as happened in the later part of the sample period.

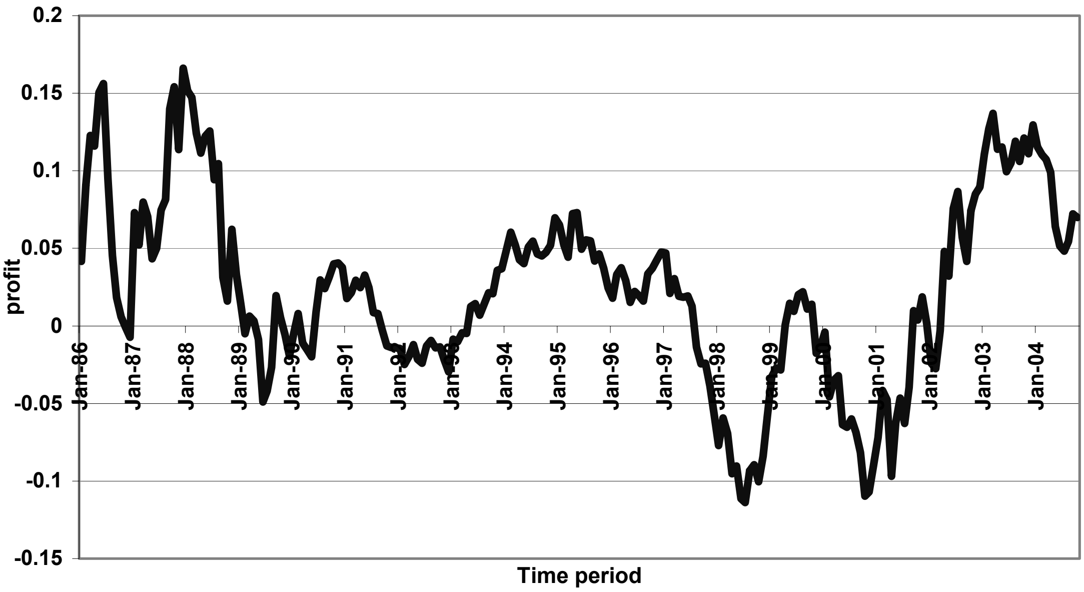

Figure 5: One-year forward profit

By the *effective exchange rate* we will mean either the actual exchange rate or the effective conversion rate ECR, depending on context. If the actual exchange rate is used, we shall refer to *R* as the natural or unhedged term of trade exposure. If the ECR is used, the result will be the hedged objective variable. In the latter case it will be convenient to use the log forward rates to accord with the desired hedge object as in section 3.1. Adapting expression (6) above, we would then have

$$logECR_{t+4} = h_1 logF_{4,t} + h_2 logF_{3,t+1} + h_3 logF_{2,t+2} + h_4 logF_{1,t+3} + (1-h_1-h_2-h_3-h_4)\times logS_{t+4}, \quad (8)$$

Note that the forward rates are hedging not only against their direct target (the spot exchange rate) but also against the CPI and product price components. The reason is that the latter variables might well be correlated with the exchange rate. This is particularly true for the product price, as the NZD is largely driven by offshore commodity prices.

The hedges considered are of the smoothing or passive type. The objective is to assist farmer production planning by smoothing out bumps and grinds, with special reference to those at the lower end. We do not consider in the present paper the issue of active or tactical hedging wherein the hedge ratio might depend upon the current state of exchange rates or commodity prices, with an implicit forecasting agenda.

3.4 The hedging decision problem

The complete optimisation model may be summarised as follows.

Notation:

S = spot exchange rate USD/NZD.

F = forward rates: for brevity F_3 = 3-month forward executed 3 months prior to spot date, F_6 means the 6-month forward executed 6 months prior to spot date, similarly F_9 and F_{12} .

P_d = price index of product.

P_s = expense price index.

R = objective variable.

The decision problem:

$$\max_{h_1,h_2,h_3,h_4} E[U(R)] = E[R-P] + \beta E[(R-P)\times SF_b(P-R)], \tag{9a}$$

where:

$$R = \log P_d + (1-h_1-h_2-h_3-h_4)\log S + h_1 \log F_3 + h_2 \log F_6 + h_3 \log F_9 + h_4 \log F_{12} - \log P_s \tag{9b}$$

$$P = VaR_{5\%}\{\log P_d + \log S - \log P_s\}; \tag{9c}$$

And either:

The options equivalent model (as in expression (5), section 2.1)

$$(R-P)\times SF_b(P-R) = [R\times N(-d_1) - P\times N(-d_2)]$$

$$\text{with } d_1 = \frac{\ln(\frac{R}{P}) + \frac{\sigma^2}{2}}{\sigma}, d_2 = d_1 - \sigma \quad ; \tag{9d}$$

Or: the logistic fuzzy model (as in Appendix B)

$$SF_b(x) = 1/(1+e^{-x/\lambda}) .$$

All subject to $h_1 \geq 0$, $h_2 \geq 0$, $h_3 \geq 0$, $h_4 \geq 0$ *and* $h_1+h_2+h_3+h_4 \leq 1$.

Commentary (equation numbering refers):

(9a): For the objective function we initially set $\beta = 8$, for reasons explained in section 2.1 (under calibration). However, we also varied this parameter to examine responses to differing degrees of risk aversion including different settings of P (see below).

(9b): Derived from equations (7) and (8). The problem is to find the optimal forward rate weights h_i .

(9c): The critical parameter P is based on the historical lower 10% VaR point for the natural exposure, i.e. the historical exchange rate S is used in place of any hedged exchange rate. In effect, we located historically the most uncomfortable zone for the farmer. Thus $P_{10} = -1.8507$, with an alternative at $P_5 = -1.8739$ as a slightly more relaxed stance (see section 3.2 above); $P_{20} = -1.8099$ is also employed.

(9d): The results reported below are based on the put option equivalent version of the fuzzy utility function. This assists with convergence of the numerical algorithm, as previously noted. We set $\sigma = 0.025$ based on the historical volatility of changes in R. The logistic fuzzy minimum with λ =0.01 was also utilised as a check, with broadly similar results.

The expected values appearing in the above maximisation problem (expression 9a) were estimated as sample averages of the corresponding magnitudes. It is not necessary for the observations on R to be *i.i.d.* over time, but it is assumed that sample values for the objective function converge as the sample size becomes large with probability one, uniformly in the hedge parameters $\boldsymbol{h}$ in a neighbourhood of the true optimum. Inspection suggests that this is true in the present context. It does not mean that log exchange rates or commodity prices are necessarily stationary.

The above formulation assumes a rolling hedge spanning a year ahead. We also investigated a longer term hedge horizon out to 3 years ahead – in this case the forwards F are for 1, 2 and 3 years.

3.5 Data and computation

Data span the period Jan 86 to Nov 04, giving 227 monthly observations. Spot and forward exchange rates are obtained from Thomson Financial Datastream as the mid rate in each case. However 1, 2 and 3 year forwards were constructed synthetically from covered interest parity, rather than being direct market data. For this purpose, interest rate swap rates were used after 1999 when they became available, while before this date swap rates were estimated by adding a credit spread to the government bond rates, validating by comparing with the post 1999 period. Zero coupon rates bootstrapped up from the swap rates were then used to compute the longer dated FX forwards. Data for long term hedging was slightly restricted to span July 87 to Nov 04. Dairy price data is derived from ANZ Bank data series used to construct their commodity price indices. The monthly farmer expense price index is interpolated using quarterly data from Statistics NZ, the official statistical agency.

The computational method used to solve the decision problem (9) is *Sqpsolvemt* from Gauss, a quadratic programming routine that employs a sequential steepest descent method. It was necessary to smooth the objective function as earlier described, and once this is done, the routine converges quickly.

IV Hedging results

4.1 One-year horizon

Table 2 shows how the optimal weights vary with changing the P and β parameters to indicate different degrees of lower tail risk aversion. Points P_5 = *-1.8739*, P_{10} = *1.8507*, and P_{20} = *-1.8099* refer to VaR points of the unhedged distribution as a way of locating discomfort points. Note that these will not necessarily be VaR points of the hedged distributions, and the latter will be given as output diagnostics. The weights h_i are formally expressed as percentages.

As anticipated, low values of the penalty parameter β amount to an absence of effective risk aversion and the recipe is to use only the one-year forward rate. As β rises, the weight given to the one-year forward diminishes, though it always remains appreciable, and a more diversified portfolio of nearer months appears. A small spot exposure is called for where risk aversion is substantial, while near months are also weighted in the case of $P = P_{10}$ which is treated as the default case in what follows.

P_5 = -1.8739 β ↓	Optimal weights *h* →	Spot	3-M F	6-M F	9-M F	12-M F
0		0	0	0	0	1
8		0	0	0	0	1
18		0	0	0	5.66%	94.34%
50		1.66%	9.02%	0	5.79%	83.53%
100		5.80%	9.11%	0	5.68%	79.41%
200		8.17%	7.38%	4.67%	3.78%	75.99%

P_{10}= -1.8507 β ↓	Optimal weight *h* →	Spot	3-M F	6-M F	9-M F	12-M F
0		0	0	0	0	1
8		0	0	0	0	1
18		0	9.28%	0	5.37%	85.35%
50		4.94%	8.94%	6.71%	4.66%	74.75%
100		7.08%	6.42%	13.25%	2.47%	70.78%
200		8.15%	5.21%	16.97%	1.18%	68.48%

P_{20} =-1.8099 β↓	Optimal weights *h* →	Spot	3-M F	6-M F	9-M F	12-M F
0		0	0	0	0	1
8		0	1.25%	2.38%	0	96.37%
18		0	10.65%	15.62%	1.99%	71.74%
50		3.06%	9.76%	23.38%	1.47%	62.33%
100		4.35%	8.89%	25.80%	1.21%	59.75%
200		5.00%	8.46%	27.05%	1.04%	58.46%

Table 2: Variation of optimal weights with tail risk aversion

Hedge performance can be depicted in several ways. Table 3 compares the optimised portfolio for β =*18* and P_{10} against simple hedges involving just one forward as well as the unhedged outcome. A simple one-year forward has a higher mean, as we should expect, but is more adverse with respect to the VaR, CVaR and expected utility of the hedged distributions.

P =P_{10} , β =18	Unhedged LOG FTT	6-m Forward fully hedged	1-Y Forward Fully hedged	Optimised portfolio
Mean	-1.6637	-1.6384	-1.6103	-1.6142
Variance	0.0245	0.0202	0.0240	0.0226
10% VaR	-1.8507	-1.8014	-1.8071	-1.8032
10% CVaR	-1.8868	-1.8328	-1.8331	-1.8282
E[U]	0.1210	0.2032	0.2297	0.2354

Table 3: Outcomes for log terms of trade

Alternatively, one can compare the optimally hedged portfolio with the unhedged, natural outcome, taking the latter as benchmark. Figure 6 is an ordered mean difference plot (Bowden 2000, 2005) choosing the optimum portfolio for β=*18* and P_{10}. The height at each value of the benchmark represents the investor or managerial surplus that results relative to the unhedged position. Lower values on the horizontal axis represent a more risk averse investor. The OMD plot is presented with one-sigma confidence bands on each side. The uniform positivity of the OMD schedule shows that when taken over the entire period, the optimised portfolio would be preferred to remaining unhedged, by any risk averse investor. On the other hand, the optimised technique does not go so far as to stochastically dominate the unhedged position.

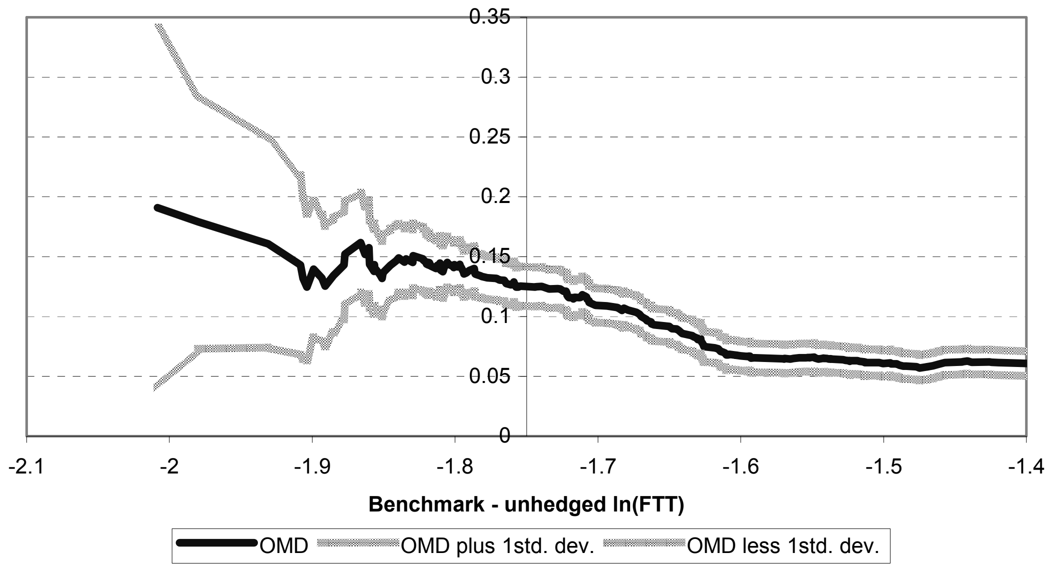

Figure 6: OMD plot for optimised portfolio against unhedged as benchmark

Superior performance over the entire period does not guarantee superiority over subintervals and this issue may be of importance at times. Figure 7 is a historical comparison of the optimised farmer terms of trade *R* (with P_{10} and $\beta = 18$) as the lighter line, and the natural or unhedged outcome as the darker. The optimised portfolio is generally better on both the downside and the upside. However, it would not have recovered as quickly after 1997-1998 as did the natural unhedged position, though the recovery in the latter was short-lived.

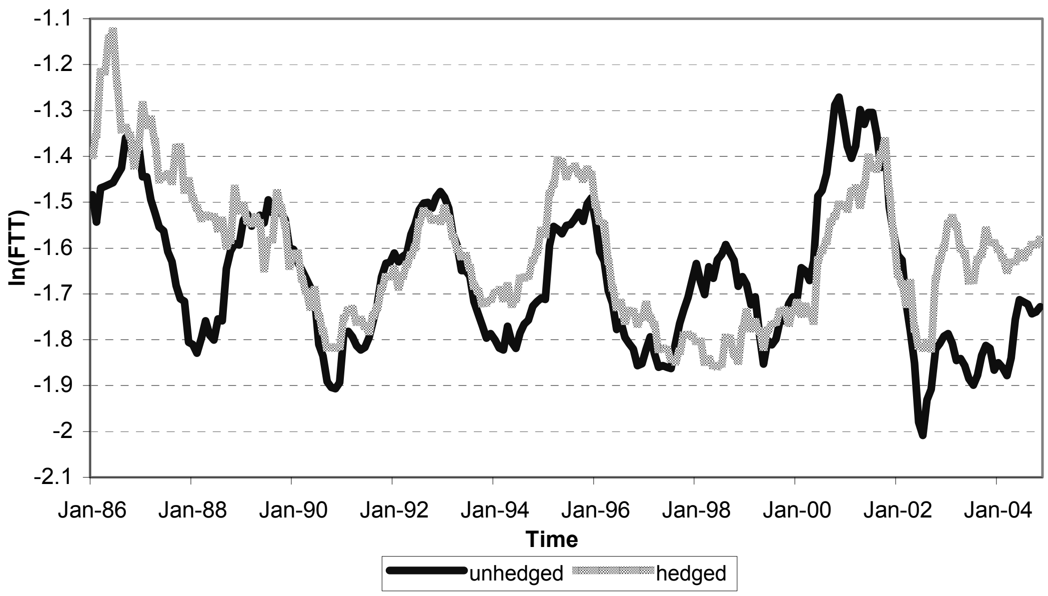

Figure 7: Historical time paths, hedged v unhedged

4.2 Three year horizon

Table 4 gives the optimal weights in the longer 1-3 year hedging context. The predominant influences come from the one-year forward, with a negligible gain from hedging three years forward. The reason lies in the correlation between the three components of the term of trade index. Table 5 reveals that, though the historical average of three-year forward rate is higher than the one year forward rate, the regression coefficient for the former against the [commodity price/expense price] is positive while that for the latter is negative. Use of the three-year forward hedged portfolio will therefore result in higher downside risk, not lower. A risk neutral exporter might still prefer three-year forwards to hedge the currency exposure, but risk averters will tend to weight more heavily the one-year forward.

P_5=-1.8739 β↓	Optimal weights h →	Spot	1-y	2-y	3-y
0		0	0	0	1
5		0	69.65%	15.39%	14.95%
8		0	74.26%	12.02%	13.72%
18		0	80.29%	6.75%	12.96%
50		7.99%	81.90%	4.16%	5.95%
100		12.46%	83.54%	0	4.01%

P_{10}=-1.8507 β↓	Optimal weights h →	Spot	1-y	2-y	3-y
0		0	0	0	1
5		0	79.47%	5.10%	15.43%
8		0	83.65%	1.60%	14.75%
18		7.34%	83.99%	0	8.66%
50		13.91%	82.41%	0	3.68%
100		15.81%	81.82%	0	2.37%

Table 4: Variation of optimal weights with tail risk aversion

	spot rate	1y forward rate	2y forward rate	3y forward rate
Historical mean	1.8097	1.8687	1.8921	1.8923
Regression coeff. on commodity price/expense price (brackets are t values)	-0.4652 (-6.8024)	-0.6846 (-10.7958)	-0.1787 (-2.2162)	0.2211 (2.4668)

Table 5: Comparisons among forward rates and spot rate

4.3 Out of sample performance

Tests of potential robustness separate the estimation period from the evaluation period. They can have different forms.

(a) *Parameter stability and its effects*

Because of the limited number of tail events, separation of the estimation period from the subsequent evaluation period does in general result in different hedge ratios, depending on the incidence of tail events in the chosen sample period. Thus if the sample period is constrained to end in May 1998, the adverse terms of trade in 2002-3 are not included and one finds that optimal hedge ratio is biased towards 100% one year forward, instead of the more distributed pattern noted in table 2. Thus it is important for the estimation phase data to include a sufficient number of tail events relative to the chosen critical or 'pain' point.

(b) *Sequential or embedded time hedges*

A second robustness procedure corresponds to that employed by Chan, Gan and McGraw (2003) to examine empirical hedging effectiveness. In order to estimate the optimized results, we should include a sufficient number of points below the pain point. Therefore, the out of sample analysis for short term hedging is started with a sample estimation period spanning Jan.1986 to Dec.1998. The estimated optimal weights from the initial sample were used for hedging against the spot exposure at Dec.1999 (in one-year hedging). This was then rolled forward one month at a time. Thus the next estimation was based on the period from Jan.1986 to Jan.1999 and the hedged exposure is at Jan.2000. At each step, a longer sample period was used for the weights estimation, the idea being to mimic the way that things might be done in practice.

Table 6 shows the diagnostics for the shorter term hedging. The optimised technique is marginally inferior on CVaR to one year forward hedging, though it has a better VaR. Both are markedly superior to the unhedged position.

Out-of-sample for 1-year *ln FTT* hedging (12/1999-11/2004)			
	Unhedged	1-Year forward hedging	Optimized hedging (P_{10} =-1.85, β =18)
Mean	-1.6633	-1.6027	-1.6109
Variance	0.0455	0.0112	0.0126
10%VaR	-1.8785	-1.7511	-1.7378
10% CVaR	-1.9353	-1.7963	-1.8098

Table 6: Sequential hedging performance comparison

Things change with to the longer 3-year horizon. The optimised and unhedged positions are now both superior to the straight 3-year forward. The optimised strategy has the smaller CVaR and is superior on the mean as well. Table 7 illustrates.

Out-of-sample for 3-year *ln FTT* hedging (12/2001-11/2004)			
	Unhedged	3-Year forward hedging	Optimized hedging (P_{10} =-1.85, β =18)
Mean	-1.8080	-1.6305	-1.6843
Variance	0.0090	0.0875	0.0198
10%VaR	-1.9036	-2.0362	-1.9168
10% CVaR	-1.9569	-2.0880	-1.9255

Table 7: Sequential hedging performance comparison

Finally, we tried powering up the bad zone as in expression (3) of section II using the fuzzy logistic smoothing as in Appendix B. As expected, a value of κ =1.5 or 2 resulted in a stronger tendency to stick with straight one-year forwards at the expense of any spot exposure.

V Concluding remarks

It is possible to devise hedging techniques and algorithms that preserve the general desirability of maximising expected returns or profit objectives, while diminishing the risk of adverse outcomes. The latter are nearly always confined to one particular zone of the outcome space, and an objective function can be calibrated to accommodate such contingencies. One can identify the risk penalty in terms of a hazard zone in which the manager has effectively written put options on the profitability outcome. The imputed value of these options is the risk penalty. The approach has a connection with value at risk and conditional value at risk. It inherits the value at risk point as a natural calibration point, as well as the focus on losses downside of this critical point. However, the explicit expected utility framework avoids some of the difficulties associated with reconciling value at risk and conditional value at risk. The smoothing implicit in options equivalence facilitates relatively simple computation routines.

Moreover, there is a correspondence with the theory of corporate finance. A rationale commonly attributed to hedging is that it helps to minimise the risk of bankruptcy, or more precisely the costs arising from bankruptcy. The existence of the latter is as though the firm has written options in favour of third party claimants to the firm's value in the event of bankruptcy (lawyers, statutory managers, liquidators, etc.). The existence of such options detracts from the wealth of equity and bond holders, who are now sharing corporate value with third party claimants. The value of these options becomes significant once lower critical points have been reached for the firm's profits or cash flow. In such an interpretation, it is these third parties who effectively hold the options that appear in the utility function.

On the other hand, the empirical limitations of this and similar methodologies should be recognised. It is desirable from the empirical point of view to have a reasonable number of observations on both sides of any chosen critical point. Calibration points such as a 5% VaR are usually going to represent rare events on the downside. Given typical data runs, this suggests that the user should adopt 10% critical points or higher, rather than rely too much on conclusions based on just a few observations in the critical zone.

Use of the corporate terms of trade for hedging decision making is useful in several ways. The economic basis is easy to understand and the hedging outcome is readily compared visually with the unhedged outcome. In addition, the bounded nature of the resulting time series enables one to draw on a variety of established statistical metrics to measure the effectiveness of hedging.

Turning to the particular context of foreign exchange and commodity hedging, our results offer qualified support for the use of the forward market where a chronic forward rate discount or premium exists. The qualifications are as follows: (a) The hedge should be distributed over maturity months or years, with a smaller spot or near-month weighting, rather than relying wholly on the forward rate. This allows for welfare loss associated with exceptional times, notwithstanding an ex ante forward discount or premium. (b) The apparent forward rate discount should not be relied on too far forward, no matter that the expected benefit might grow in forward maturity time. Movements in commodity prices, as another exposure element, may destroy the gains altogether. Asymmetric hedging methodology applied to the corporate terms of trade assesses the potential effect of such correlations, balancing up the sources of expected gain and contingent loss. It is worth emphasising that the use of forward markets is not necessarily a response to risk aversion. Indeed, our results suggest that the primary motive is simple expected return and the component of unhedged spot, where it appears, arises out of increasing risk aversion. Thorp (2005) makes a similar point in the context of funds management.

Finally, the empirical work has been directed to the issue of the forward premium and its usage. It does not address the further issue of whether options should be used, either on a stand-alone basis or in conjunction with forwards. On the face of it, like should generate like, so that if the manager is in the natural position of effectively having to write a put option (as in CVaR), then this could be fixed by buying an offsetting put option in the adverse zone. In effect, one would be using options to hedge a natural option-type position. Our results suggest that given the gains to be made from using forwards, use of options may be an expensive supplement, given the volatility of the NZ dollar in the chosen context. However, the optimal combination of forwards and options in general remains a subject for further research.

Appendix A

The generalised Rubinstein risk premium for the segmented utility function

The GR risk premium is defined for an arbitrary risk-averse utility function *U(R)* by

$$\theta = -\frac{E[(R-\mu_R)U'(R)]}{E[U'(R)]}.$$

It uses the marginal utility weights to adjust the expected outcomes; the weights have a close correspondence with the state price deflators used in martingale pricing (see Bowden 2005). The meaning of the risk premium is that if a return or object of certain value μ_R - θ was available, then the investor or manager would be indifferent if the last investment dollar or unit was devoted either to *R* or to the certain asset. Adapting the utility function (1) to the present context, we have

$$\theta = \frac{\beta F_R(P)}{1+\beta F_R(P)}[\mu_R - \mu_R(P)], \tag{A1}$$

where $\mu_R = E[R]$ and $\mu_R(P)$ is defined as the censored mean $E[R|R \le P]$. Note that $\mu_R(P) < \mu_R$, for any *P*, so the risk premium is always positive and in increasing with β.

To prove formula A1, it is easiest to use methods from generalised functions (see Bowden 2005, Appendix C) or else Lighthill (1959), a standard reference. The same result can be obtained – at greater length - by breaking up the domains of integration and using more traditional methods. The utility function is

$$U(R;P) = R - P + \beta(R-P)SF(P-R).$$

Also $SF(P-R) = 1 - SF(R-P)$ and $\frac{d}{dR}SF(R-P) = \delta(R-P)$, the Dirac delta function. In addition, for any smooth function ϕ(R),

$$\int_{-\infty}^{\infty} \phi(R)\delta(R-P)f(R)dR = \phi(P)f(P).$$

Hence $E[\phi(R)\delta(R-P)] = \phi(P)f_R(P)$.

Finally, $E[SF(P-R)] = F_R(P)$ and $E[RSF(P-R)] = F_R(P)\mu_R(P)$ by definition of the censored mean $\mu_R(P)$. The desired result follows by substitution.

Appendix B

Using fuzzy logic to smooth out the kink

For a continuous random variable, the step function *SF(x)* is an indicator function for two distinct sets of positive measure, namely the positive and negative numbers, so that a given number *x* belongs in either one or the other. A standard fuzzy indicator function would attach a number between zero and unity to indicate the strength of the classification. Thus one could approximate the step function in expression (1) of the text by a Normal distribution function and write $SF(x) \sim N(x;0,\sigma^2)$ for a suitably chosen value of σ. The Logistic distribution function is also useful: $SF(x) \sim 1/(1+e^{-x/\lambda})$. By setting the 'smearing' parameters σ or λ arbitrarily small, one can approach closer and closer to the 'all or nothing' switch given by the unit step function. Figure B1 depicts a fuzzy generalized value at risk (GVaR) utility function, using the logistic version, which is marginally faster to compute. A value of $\lambda = 0.01$ suffices for a fairly close approximation and yields much improved convergence in the empirically based optimisation reported below. For other uses of fuzzy logic in Finance, see Simonelli (2001), Tseng (2001), Zmeskal (2001, 2005).

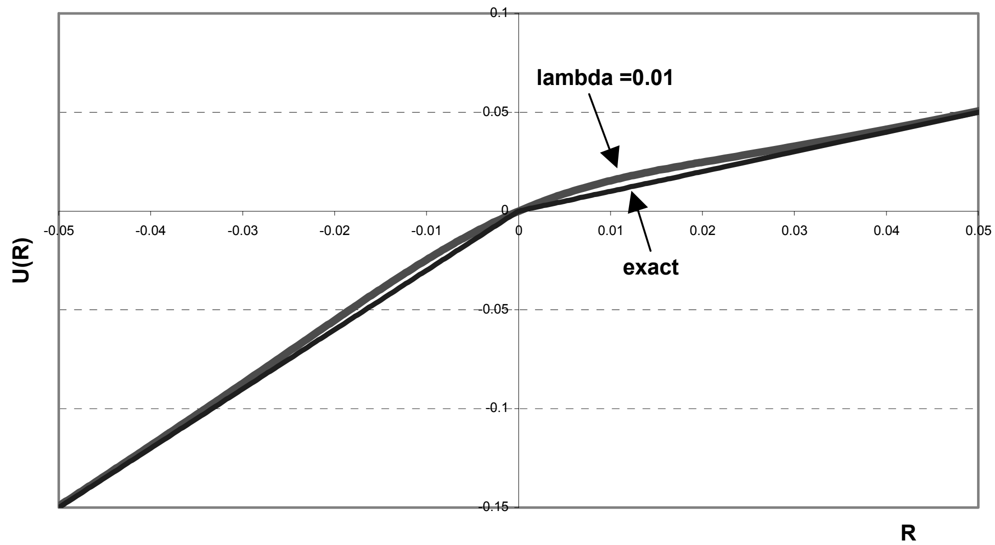

Figure B1: Fuzzy approaches to the exact utility function (P =0)

Comparing figure B1 and figure 2 of the text for the option equivalent method, it will be evident the fuzzy method smoothes from above whereas the option equivalent method smoothes from below. One might therefore expect the option equivalent utility function to be marginally more sensitive to behaviour approaching the value at risk point *P* from above, and hence slightly more defensive, which turns out to be the case. In all cases, the intent is to preserve the sharp curvature at the VaR point and the aggravated penalty slope to the left of this point.

Appendix C

Mark to market v. ECR hedging: outcome measurement

Consider the four period rolling hedge of section 2.2. Since the forwards are marked to market at the end of each period, the forward can be considered as rolling over continuously. Suppose the manager sells US$25m 1-year forwards on forward rate $F_{4,t}$ at time t. At time $t+1$, he buys back US$25m 1-year forwards and sells US$25m 9-month forwards on $F_{3,t+1}$. At time $t+2$, he buys back US$25m 9-month forwards and sells US$25m 6-month forwards at $F_{2,t+2}$. At time $t+3$, the manager buys back US$25m 6-month forwards and sells US$25m 3-month forwards at $F_{1,t+3}$. At time $t+4$, he buys back $25m 3-month forwards.

The figure on the facing page illustrates the resulting cash flows in each period as they accrue to changes in values of forward contracts, making allowance also for any spot transactions. Real time is measured vertically and maturity time is horizontal. We evaluate the portfolio in the same way as we construct it.

(a) Under the *effective conversion rate approach*, the portfolio is valued by adding up all the cash flows resulting from changes in forward rate contracts that all refer to the same specific time. See the example highlighted with grey colour.

Return attributable to time t+4 =

$$25 \times (F_{4,t}-F_{3,t+1}) + 50 \times (F_{3,t+1}-F_{2,t+2}) + 75 \times (F_{2,t+2}-F_{1,t+3}) + 100 \times (F_{1,t+3}-S_{t+4}).$$

Combining this equation with the spot position, we get

$$25 \times (F_{4,t}-F_{3,t+1}) + 50 \times (F_{3,t+1}-F_{2,t+2}) + 75 \times (F_{2,t+2}-F_{1,t+3}) + 100 \times (F_{1,t+3}-S_{t+4}) + 100 \times S_{t+4}$$

$$= 100 \times (0.25 \times F_{4,t} + 0.25 \times F_{3,t+1} + 0.25 \times F_{2,t+2} + 0.25 \times F_{1,t+3}).$$

Since the amount of exposure is $100, the part in the bracket is the effective conversion rate. More generally, let h_1, h_2, h_3, and h_4 be the weights for one year forward purchased at time t, the nine month forward bought at time $t+1$, six month forward acquired at time $t+2$ and three month forward entered into at time $t+3$. The effective conversion rate will become:

$$h_1 \times F_{4,t} + h_2 \times F_{3,t+1} + h_3 \times F_{2,t+2} + h_4 \times F_{1,t+3} \text{ ; with } h_1+h_2+h_3+h_4=1 .$$

(b) The *cash flow approach* disregards the relation existing among forwards and the corresponding exposures. Instead, the focus is only when the cash flow happens. The value of the hedging portfolio is indicated by the sum of all the cash flows occurring at the current period. An example for real time *t+4* is specified with blue colour in the above table.

Cash flow at time t+4 =

$$100\times(F_{1,t+3}-S_{t+4})+75\times(F_{2,t+3}-F_{1,t+4})+50\times(F_{3,t+3}-F_{2,t+4})+25\times(F_{4,t+3}-F_{3,t+4})$$

Comparison

The effective conversion rate approach build up the hedging portfolio in line with the underlying exposures. As a given currency exposure is supposed to be hedged by a series of contracts, several exchange rates exist in the hedging portfolio at the time of evaluation. An effective conversion rate by averaging the assorted exchange rates alleviates the measurement complexity. However, the approach ignores the time value of money and simply adds up the cash flows occurring at different time periods. Moreover, under the effective conversion rate approach, hedgers only evaluate the portfolio when the exposures mature, without regarding the risk of cash outflow before maturity.

The cash flow approach encompasses the time difference of cash flows, which is missed by the effective rate approach. This cash flow approach defines the portfolio as a mixture of contracts existing at the time of evaluation and values the cash flows when they take place. It is therefore more financially correct as an indication of ongoing changes in net worth. However, such an approach fails to establish a direct correspondence between exposures and hedging instruments, and this is especially true when the given spot exposures vary in quantity over time.

Time	t+1	t+2	t+3	t+4	t+5	t+6	t+7	......
......								
t+1	$25\times(F_{1,t}-S_{t+1})$	$25\times(F_{2,t}-F_{1,t+1})$	$25\times(F_{3,t}-F_{2,t+1})$	$25\times(F_{4,t}-F_{3,t+1})$				
t+2		$50\times(F_{1,t+1}-S_{t+2})$	$50\times(F_{2,t+1}-F_{1,t+2})$	$50\times(F_{3,t+1}-F_{2,t+2})$	$25\times(F_{4,t+1}-F_{3,t+2})$			
t+3			$75\times(F_{1,t+2}-S_{t+3})$	$75\times(F_{2,t+2}-F_{1,t+3})$	$50\times(F_{3,t+2}-F_{2,t+3})$	$25\times(F_{4,t+2}-F_{3,t+3})$		
t+4				$100\times(F_{1,t+3}-S_{t+4})$	$75\times(F_{2,t+3}-F_{1,t+4})$	$50\times(F_{3,t+3}-F_{2,t+4})$	$25\times(F_{4,t+3}-F_{3,t+4})$	
.....								

Table C1: Cash flows from rolling hedges

References

Alexander, G. J. and A.M. Baptista (2004) A comparison of VaR and CVaR constraints on portfolio selection with the mean-variance model. *Management Science*, 50: 1261-1273.

Artzner, P., F. Delbaen, J. M. Eber and D. Heath. (1997) Thinking coherently. *Risk*, 10, 11, 68-71.

Artzner, P., F. Delbaen, J. M. Eber and D. Heath (1999) Coherent measures of risk. *Mathematical Finance*, 9, 203-228.

Basak, S. and A. Shapiro (2001). Value-at-risk-based Risk Management: Optimal Policies and Asset Prices. *Review of Financial Studies*, 14, 371-405.

Bowden, R. J. (2000) The ordered mean difference as a portfolio performance measure. *Journal of Empirical Finance*, 7, 195-223.

Bowden, R. J. (2006) The generalised value at risk admissable set: constraint consistency and portfolio outcomes, *Quantitative Finance*, forthcoming.

Bowden, R.J. (2005) Ordered mean difference benchmarking, utility generators, and capital market equilibrium, *Journal of Business*, 78, 441-467.

Bowden, R.J (2005) *Macroeconomics and Financial Management*. Wellington: Kiwicap Education.

Chan, K. F., C. Gan, and P. A. McGraw (2003) A hedging strategy for New Zealand's exporters in transaction exposure to currency risk. *Multinational Finance Journal*, 7(1/2), 25.

Ferreira, Miguel A. and Carlos Gonçalves (2004) Tail Risk and pK-Tail Risk. Working paper. ISCTE Business School- Lisbon. http://ssrn.com/abstract=639181

Jorion, P. (2000) *Value at Risk: The New Benchmark for Managing Financial Risk*, McGraw-Hill.

JP Morgan Ltd (2001) *Return to RiskMetrics – the Evolution of a Standard*.

Lighthill, M. J. (1958) *An introduction to Fourier analysis and generalised functions.* New York: Cambridge University Press.

Palmquist, J., P. Krokhmal and S. Uryasev (2001) Portfolio optimization with conditional value-at-risk objective and constraints. *Journal of Risk*, 4, 2.

Rockafellar, R. and S. Uryasev (2002) Conditional value-at-risk for general loss distributions. *Journal of Banking & Finance*, 26, 1443-1471.

Simonelli, M. R. (2001) Fuzziness in valuing financial instruments by certainty equivalents. *European Journal of Operational Research*, 135, 296-302.

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

Tseng, F.-M., G.-H. Tzeng, H.-Ch. Yu, and J.C. Yuan (2001) Fuzzy ARIMA model for forecasting the foreign exchange market. *Fuzzy sets and systems*, 118, 9-19.

Uryasev, S. (2000) Conditional Value-at-Risk: Optimizing Algorithms and Applications. *Financial Engineering News*, 14.

Yamai, Y. and T. Yoshiba (2002) Comparative analyses of expected shortfall and value-at-risk (2): expected utility maximization and tail risk. *Monetary and Economic Studies.*

Zadeh, L. A. (1965) Fuzzy sets. *Information control*, 8, 3, 338-353.

Zadeh, L. A. and R. Bellman (1970) Decision-making in a fuzzy environment. *Management Science*, 17, 141-164.

Zmeskal, Z. (2001) Application of the fuzzy-stochastic methodology to appraising the firm value as a European call option. *European Journal of Operational Research*, 135, 2, 303-310.

Zmeskal, Z. (2005) Value at risk methodology under soft conditions approach (fuzzy-stochastic approach). *European Journal of Operational Research*, 161, 337-347.