February 2003

Utility-Invariant Linear Hedges

by Roger J. Bowden[1]

Abstract

Empirical hedge procedures typically adopt a symmetrical mean square error immunization criterion, wherein positive values of the hedging basis or net hedged return are penalized equally along with negative. The econometric hedge, which proceeds by estimating a presumed underlying regression relationship, is naturally adapted to this loss function. In many circumstances, however, a symmetrical loss function may be excessively limiting or even inappropriate. This paper presents conditions under which a given hedge such as the econometric hedge is unaffected by the choice of utility function. If an invariant hedge exists, it can be consistently estimated by OLS, but the converse is not true: classical regression conditions do not suffice for the hedge to be invariant. Ordered mean difference techniques can be used to adjust the econometric hedge up or down according to the user's perception of sensitive zones of net return, as in value at risk.

Key Words: Econometric hedge, futures, hedging, inflexion point, invariant linear hedge, ordered mean difference, portfolio enhancement, stochastic volatility, value at risk, utility generator.

---

[1] The author is with the School of Economics and Finance, Victoria University of Wellington, P.O. Box 600 Wellington New Zealand, email roger.Bowden@vuw.ac.nz or roger.Bowden@kiwicap.co.nz. He is grateful to Richard Harris and Jian Shen for data and discussion on the topic.

I Introduction

Most corporates hedge their stochastic exposures at some time or other. Yet hedging as a subject is controversial on a theoretical level and problematic in its empirics. On a theoretical level, doubts have centred upon whether the corporate should do the hedging or whether they should leave it to their shareholders to achieve their own desired risk exposure. This amounts to the position that the net market value of the hedge (e.g. a forward position) is zero, so that the value of the hedged firm should remain just the same as the unhedged. Contrary positions invoke explanations such as controlling the exposure to bankruptcy costs; the ability that hedging confers to take advantage of unforeseen investment opportunities; or the non alignment of risk preferences of managers with those of shareholders. The empirical debate inherits some of the theoretical divisions. For example, most of the extensive econometric literature takes as given the immunization criterion, wherein the loss function is explicitly or implicitly assumed to be the mean square error. This penalizes equally departures in either direction, so that a metals producer with natural long position in spot prices would be protected via a short position in futures if the spot price fell, but would forego gains if the spot price did in fact rise. Immunisation would appear to be linked most naturally with the bankruptcy argument. But if this were so, one could legitimately ask whether this extra degree of risk aversion should figure in the design of the hedge, so that perhaps the conventional econometric hedge would not provide sufficient protection.

On the other hand, the econometric literature has helped us to focus on the potential importance of non classical regression assumptions. Thus assuming that an underlying structural model exists, the disturbance terms are often non Gaussian (e.g. Bolleslev, Engle and Woolridge 1988, Baillie and DeGennaro 1990); there may be GARCH type dependence (e.g. Baillie and Myers 1991, Kroner and Sultan 1993, Lin *et al* 1994, Brooks and Chong 2001, Park and Switzer 1995); Gagnon and Lypny 1995); or there may be leptokurtosis in any of all of the spot price, the hedge instrument, or the residual connecting them (e.g. Germat and Harris 2002, Harris and Shen 2002).

However, the underlying decision framework remains the same. The hedge ratio is estimated as the structural parameter(s) connecting the spot and the hedge instrument, effectively reducing the hedge calculation to an econometric estimation problem. The two mesh quite nicely: the statistical loss functions that underly econometric estimation are fundamentally the same as those that are presumed (in default) to underpin the reason for hedging to be carried out. The result has therefore been long in econometric virtuosity, but uneasily short on the economic reasoning as to why the hedge is being carried out and how this might affect things. Some parallels exist with the derivatives literature, which again is long on the technical specifics, but short on just why and when they should be used.

Does it matter? It might be the case that the econometric hedge is robust to a wide variety of underlying hedging preference functions. The leading case is where the spot and hedge instrument returns are jointly Gaussian. It then turns out (see section II) that no matter what the precise preference function, the optimal hedge ratio is just the econometric beta (e.g. OLS, where appropriate). In other words, suppose my preference function is now asymmetric, so that I dislike negative residual returns but like positive returns, as distinct from the immunization preference, which penalizes positive residual returns, as does the econometric loss function. Provided those returns are normally distributed, the non matching does not matter. I should still use the econometric beta as the hedge. In the language of the present paper, we should say that the econometric beta is invariant to the choice of utility function , or just 'invariant', where the context is clear. Although a bit surprising at first sight, this sort of result should be quite familiar from portfolio theory. There, no matter what the preferences, whether quadratic or not, if returns are Gaussian, mean variance applies (with minor extensions to elliptical distributions). The connection is not accidental, for the hedging problem is essentially a portfolio problem. It differs only in that adding hedges such as forwards or futures theoretically does not entail extra capital, so this is zero capital 'enhancement' or 'overlay' in the usage of the trade literature.

As earlier remarked, however, returns are not typically Gaussian. Hence we need to examine the robustness of the hedge ratio to different assumptions about the nature of preferences, or agent utility. For example if the agent's utility function encompasses a value at risk boundary for low values of the net outcome, then it may be optimal to over-hedge relative to the econometric hedge. Or the agent's preference might be state dependent, so that the degree of risk aversion would rise in bad times but fall in good times. Examples might include portfolio insurance, which varies according to the stock market cycle, or an impending danger of liquidity problems for financial institutions.

The present paper attacks the invariance problem from a variety of directions. Section II develops a general criterion for the existence of invariance, based on the joint densities of the spot and hedge instruments or their relationships. A linear hedge will be invariant if the conditional expectation of the hedge instrument returns is independent of the basis, i.e. the residual, or 'expectationally invariant with respect to the residual'. The condition does not correspond to the classical regression specification, which says that the conditional expectation of the residual (here, the basis) must be independent of the hedge instrument. Thus underlying regression relationships do not imply that the estimated hedge ratio is best or invariant, contrary to popular supposition. On the other hand, the invariance condition is indeed linked to regression consistency, as it implies that regressor and disturbance are uncorrelated, though the converse is not true. Thus *if an invariant hedge does exist*, it can always be consistently (perhaps not efficiently) estimated by least squares.

In general, symmetry of one kind or another tends to favour the existence of an invariant hedge. Even where this is not the case, it turns out that the econometric hedge is 'special', in that it can entail dead zones, such that the optimal hedge does not depart much from it as risk aversion varies. Mathematically, there is a point of inflexion involved, though this can depend upon the precise nature of the underlying structural model connecting the spot and hedge instrument returns. We develop a framework to represent such effects, drawing for this purpose upon the representations of ordered mean difference (OMD) technology (Bowden 2000, 2002, 2003).

Empirically, one can use OMD technology to diagnose whether linear invariance holds, based on the given historical data, and to suggest appropriate adjustments where it does not. In local invariance, one sets up the econometric hedge as the provisional position, and the OMD schedules then tell us whether we should over or under hedge relative to the beta hedge. Using the OMD utility generators in conjunction with sensitive return zones enables a closer approach to finding globally optimal linear hedges.

A crucial feature of an effective hedge is that it should be effective at dislocative or even disastrous times; it is much less important that it work well when times are good.

Econometric tests of goodness of fit or estimator methodology often adopt an averaging over all areas of the return axis that obscures this desired empirical asymmetry. If the overall fit is acceptable, but it is bad when spot returns are low, the hedge has failed in an essential function. By virtue of the orderings in its construction, the empirical OMD methodology pays closer attention to hedging basis departures than does the traditional econometric hedge.

The scheme of the paper is as follows. Section II develops the invariance theory. After briefly demonstrating invariance for the special case of Gaussian returns, we develop a window into the general invariance issue which derives from OMD theory, This comes in both local and global versions, and is then used to derive the principal theoretical result on the existence of linear invariant hedges. Section III discusses these conditions in terms of some standard models of volatility, such as ARCH and stochastic volatility, showing the importance of symmetry in return distributions, and how this relates to issues of sensitive zones of loss. Section IV has an empirical illustration based on hedging the FTSE 100 index with SPI futures. Section VI has some concluding remarks.

II Invariance

Let s be the return to be hedged (for example, the logarithmic change in a spot price), and let the hedging return be denote by f (e.g. the corresponding futures return). The return to the hedged position will be s - hf, where h is the hedge ratio. This can be regarded either as a portfolio return (for which we use R) or as a basis spread or residual (for which we shall sometimes use the symbol u); the symbols R or u (lower case) will thus be used interchangeably for the same magnitude, depending upon context.

In the immunization approach to hedging, the objective is to minimize some distance measure of R from zero. The commonly used least squares metric is $E\,[R^2]$, which results in the optimal hedge $h = E[s|f]$, conventionally estimated by least squares or econometric adaptations to allow for specific distributional assumptions. Characteristic of such ‘statistical’ approaches is their directional indifference: positive values of R are penalized equally with negative values. However, agent preferences are often more naturally asymmetric. For example, a metals producer (long in return s) would like a positive value of R but would dislike a negative value. Because he is sufficiently risk averse, so the immunization story goes, he is willing to forgo the prospect of positive departures in order to enjoy the benefit of protection on the downside.

A compound return of the form R = s –hf may be regarded as the return on a portfolio long in s and short in f. In a pure hedging context, the element f is a portfolio enhancement, requiring no additional capital (as in say a forward or a future). In any case we could imagine an increasing, concave (risk averse) Von Neumann-Morgenstern utility function U(R) such that h is to be chosen to maximize E[U(R)] , where the expectation is conditional upon current information, and taken over the joint distribution of s and f. The least squares criterion (essentially $U(R) = -R^2$) does not qualify in these terms, as it is not monotonically increasing in R; this is a formalisation of the critique mentioned above. In other words, if S_U indicates the allowable set of Von Neumann-Morgenstern utility functions, then the least squares criterion does not belong to this set. The latter is a much more suitable preference class for any discussion of risk management and will form the basis for the invariance theory of the present paper. A given hedge will be invariant if it is the same and optimal for all members of S_U .

If a hedge ratio h exists that is linear and invariant to any choice of the utility function $U \in S_U$, we shall describe it as an *invariant linear hedge*. This would be a very convenient state of affairs. The first result that follows shows that if the returns s,f are jointly Gaussian, and have zero means, then the optimal hedge ratio is indeed invariant, and moreover coincides with the least squares hedge. This sort of result is well known from mean variance analysis, hence the treatment will be brief.

The Gaussian case

Setting up the optimand as $\max_h E[U(s-hf)]$, the first order condition for h is given by

(1) $E[fU'(s-hf)] = 0$.

Applying Price's Lemma[2], we obtain the optimal hedge ratio as

[2] If x,y are Gaussian, then

$$E[g(x)h(y)] = E[g(x)]E[h(y)] + \frac{1}{2!}\sigma_{xy}E[g'(x)]E[h'(y)] + \frac{1}{3!}\sigma_{xy}^2 E[g''(x)]E[h''(y)] + \ldots$$

For more detail see Bowden (1997), or for special cases, Stein (1973) and Amemiya (1982).

(2) $$h = \frac{1}{\sigma_f^2}\left(2\mu_f \frac{E[U']}{E[U'']} + \sigma_{sf}\right),$$

where the utility derivatives are evaluated at R = (s - hf); μ_f and σ_f are the mean and standard deviation of f, and σ_{sf} is the covariance between f and s. The first order condition (1) will be a unique maximum for h if U is strictly concave.

If $\mu_f = 0$, as we shall henceforth assume, then the optimal hedge will evidently be independent of U, and will be identical to the regression coefficient of s upon f. This amounts to setting up a formal regression model

(3) $$s = \beta f + u,$$

where the residual u has zero mean and is statistically independent of the 'regressor' f. Given the assumption that s,f are jointly Normal, then a linear regression relationship automatically applies. The optimal hedge is $h = \beta = \sigma_{sf} / \sigma_f^2$.

The non Gaussian case

In practice, returns distributions are not usually Gaussian, even conditionally given the current information. This may be a matter of fat tails in the marginal distributions of f and s. Or it may be that the distribution of the disturbance term in (3) has special properties: at times of financial stress, the return on the edge instrument may depart from the usual pricing basis, e.g. from liquidity squeeze or panic motives (see section IV for an example from 1987). The latter case we shall sometimes refer to as a *basis dislocation* of f and s.

Where the optimal hedge ratio is not invariant, the choice of utility function will matter. In practice, the hedger may not even know the exact nature of the utility function, or it may not be temporally stable, depending upon the firm's particular circumstances at the time. Hence the optimal choice is more likely to be governed by an appreciation of range behaviour: certain outcome zones may be viewed as particularly sensitive. Or we may simply be interested in the way that the hedge ratio depends upon increasing or decreasing risk aversion.

Even in such cases it is often convenient to begin by postulating a linear econometric model of the form (3). One then estimates β as efficiently as possible, under various assumptions about the residual process (ε) or the distributions of s and f . The resulting estimate is used as the hedge ratio: $h = \hat{\beta}$. We shall refer to this as the *econometric hedge* or the *beta hedge*. It is often a good starting point in the determination of hedge ratio invariance.

Invariance and ordered mean difference technology

A first approach to the invariance issue is incremental in nature. One starts with a given 'provisional' or 'base' hedge ratio h, such as the regression hedge, and asks whether it is possible to improve on this in one direction or the other. Ordered mean difference techniques (Bowden 2001,2002, 2003) are a convenient way to accomplish this. OMD methods were originally developed within a framework of conventional portfolio theory, in which portfolio proportions add to unity. In the present application we are likewise proposing to add a new portfolio asset, but it requires zero capital, as for example with a futures or forward contract. This is often referred to in the informal professional literature as a portfolio 'enhancement', or 'overlay'. Sometimes this has a connotation that the risk management or enhancement process is to be undertaken by a third party, as is commonly the case in foreign exchange hedging.

The zero capital aspect notwithstanding, the reasoning is closely similar. Directional methods indicate the sign of any departure, i.e. whether we should contemplate increasing or decreasing the hedge ratio relative to the beta hedge. Given an existing or proposed hedge ratio h, we ask what an extra unit of the hedge instrument would be worth in return terms, measuring the latter as the implied tax or penalty (τ) we would be willing to pay before reverting to the provisional or base position. If the answer is zero, then h is indeed the optimal hedge. If $\tau>0$, this is a signal that we would have to be paid to give up the extra unit of the hedge, so that the given amount h is suboptimal; we need to increase it; conversely if $\tau<0$. One could think in terms of accepting or rejecting a working hypothesis that the econometric beta (for instance) is the optimal hedge. A suggestion may have arisen that the hedge should be decreased relative to the beta hedge, but if directional indicator (τ) points to 'increase', then one should be inclined to reject the suggestion.

Thus let h be the proposed hedge ratio, so the base hedge position is R = s-hf. Now add an additional amount δ of hedge in the same direction, but we shall penalize its return with a notional tax τ, so that the increment has net return f-τ. Thus the total portfolio return $R = s - h f - \delta\ (f\text{-}\tau)$. We now ask what the 'tax' τ has to be in order to drive the desired incremental holding δ down to zero. Technically, τ will solve the optimizing problem:

$$Arg\ [max_{\delta}\ E_{s,f}[U(s\text{-}hf\text{-}\delta(f\text{-}\tau))]] = 0.$$

The solution is

(4) $$\tau_u = \frac{E[fU'(s-hf)]}{E[U'(s-hf)]}.$$

If the optimal hedge exceeds the given hedge h, then τ will be positive; if it is less than h, then τ will be negative.

If the directional indicator τ is zero regardless of the particular utility function U, then we could say that the proposed hedge ratio h is *best locally invariant* among all linear hedges. As it stands, formula (4) is specific to the particular utility function U, so to check for local invariance, one would have to check for all possible utility functions. Obviously this cannot be done directly, but it can be done indirectly, by using the device of OMD utility generators.

The OMD generator utility functions constitute a family of the form

(5) $U_P(R) = \min(R - P, 0)$,

indexed by parameter P which has the dimension of the return R. Figure 1 illustrates. The point P is the 'node' or focus, and as P increases from P to P', the utility function $U_P(R)$ becomes less risk averse as a function of R. Intuitively, the utility function becomes more and more linear relative to a fixed return distribution for R.

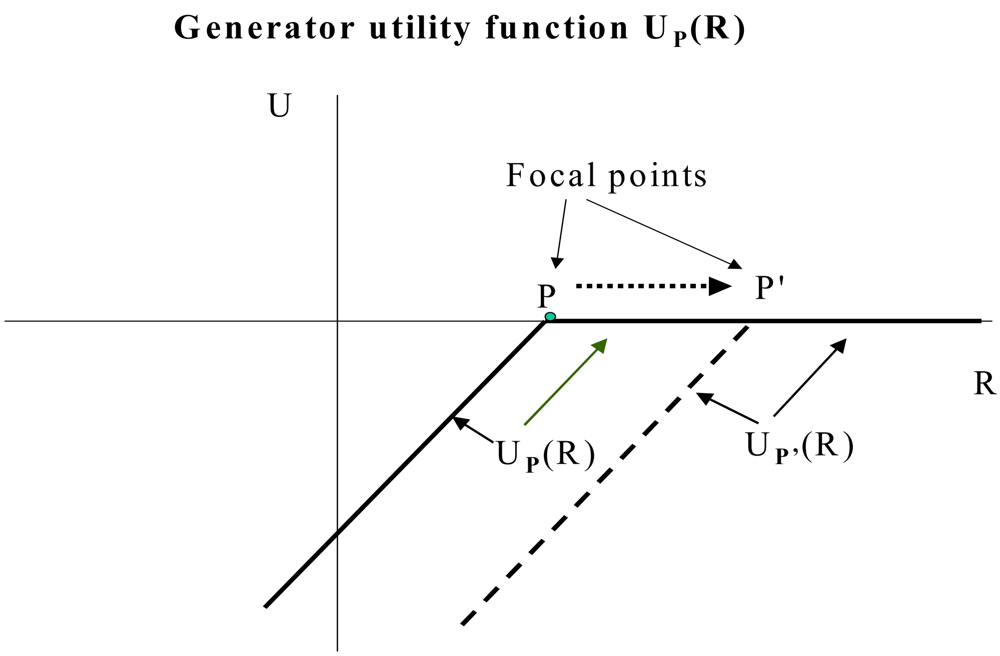

Figure 1: Generator utility functions

Applying formula (4) to $U(R) = U_P(R)$, we obtain[3]

(6) $$\tau(P) = \frac{E[fU_P'(s-hf)]}{E[U_P'(s-hf]}.$$

Expression (6) is of the nature of a truncated or censored mean, where the censoring is applied according to values of R = s-hf less than or equal to the given return number P. Writing $R_i = s_i - hf_i$, the sample counterpart can be taken as:

(7) $$\tau(P) = \frac{1}{N(R_i \le P)} \sum_{R_i \le P} f_i \quad ,$$

where $N(R_i \le P)$ means the count of R_i that are less than or equal to P. The generator equivalent margin $\tau(P)$ can be recovered by first arranging the observations in ascending order of R_i, then taking the mean of the f_i over those observations for which the associated R_i are less than or equal to the given number P. The operation can be performed on an Excel spreadsheet using the *tools/sort* menu facility. We shall refer to

[3] The derivative in (A1) is defined in the generalized distributional sense (e.g. Lighthill 1959, Antosil at al 1973). Standard differentiation rules continue to apply.

the function τ(P) as the *OMD hedge schedule,* and section IV gives some empirical examples.

Suppose that an otherwise arbitrary utility function U is increasing in R, twice differentiable and has marginal utility U'(R)→0 as R→∞ . It can be shown (Bowden 2000) that the associated equivalent margin τ_u is a positive weighted combination of the OMD generator values τ(P) :

$$(8) \qquad \tau_u = \int_{-\infty}^{\infty} w(P)\tau(P)dP \; ; \; w(P) = -\frac{U''(P)}{E[U'(R)]}F(P) \quad .$$

In formula (8) the range of $R = s - hf$ has been taken as infinite, and F refers to the distribution function of R .

Expression (8) gives us a way of testing for local invariance, relative to a given provisional h such as the beta hedge. Suppose we compute the OMD hedge schedules and find that τ(P) = 0, all P. It follows immediately that the equivalent margin τ_u must be zero regardless of the particular utility function U.

Global invariance

The equivalent margin refers to local variations about the presumed hedge h. However a more global approach can also be devised. The spanning property of the OMD utility generators extends to the utility functions themselves, as well as the equivalent margins. If U(R) is 'smooth', meaning that is has derivatives of all orders, then for purposes of expected utility maximization we can effectively write

$$(9) \qquad U(R) = \alpha_0 + \alpha_1 \int_{-\infty}^{\infty} (-U''(P))U_P(R)dP \ ,$$

where α_0 and α_1 (>0) are arbitrary constants. The spanning result (9) is implicit in the classic Rothschild~Stiglitz (1970) result that all risk averters will prefer a distribution second order stochastic dominant over another; see Bowden (2003) for a more formal proof in a portfolio context. Given expression (9), one can show[4] that if h* is the optimal

---

[4] Use expression (9), followed by the second mean value theorem of integral calculus, to write the first order condition for h as

$$E[fU'(s-hf)] \propto \int_{-\infty}^{\infty} (-U''(P))E[fU_P'(s-hf)dP = E[fU_{P*}'(s-hf] \int_{-\infty}^{\infty} (-U''(P)dP \quad ,$$

for some number P* . For this to equal zero, the first order condition applicable to U_{P*} must hold.

hedge solution for U, then there must exist some P* such that h* is optimal for the generator $U_{P*}(R)$. In other words, suppose we solve the hedge ratio decision problem for the generator utility function $U_P(R)$ and let the result be h(P), collectively amounting to a schedule of generator hedge ratios over different values of P. Then the hedge ratio for any otherwise arbitrary smooth utility function must plot somewhere along this hedge ratio schedule.

Thus if the schedule h(P) turns out to be flat, i.e. not varying with P, it will follow that the optimal hedge ratios will be the same for any smooth utility function. Evidently this is a sufficient condition for global invariance.

Theoretical conditions for invariance

It is not necessary to have Gaussian distributions, or even an underlying regression model, to ensure hedge invariance. The following result gives necessary and sufficient conditions for local invariance.

Theorem 1: Suppose that returns s and f have zero means. Then the optimal linear hedge ratio h is locally independent of the choice of utility function if and only if there exists a decomposition

(10) $s = hf + u$,

such that the following equivalent statements hold:

(i) $E[f \mid u] = 0$ *almost everywhere in u ;*

(ii) $E[f\,\delta(u-u^0)] = 0$ *for all numbers* u^0 *, where* $\delta(\cdot)$ *is the Dirac delta function.*

Proof Sufficiency follows more or less immediately from the iterated expectation applied to the first order condition for an arbitrary utility function U. Write this as

$$0 = E[fU'(s-hf)] = E[fU'(u)] = E_u[U'(u)E_f[f|u]] = 0.$$

A generator based argument allows necessity as well as sufficiency. Invariance will hold if and only if for every number P, $\tau(P)=0$. From the first order condition (4), and its non zero denominator, this will be true if and only if

(11) $E[fU'_P(s-hf] = 0$.

Indeed we can write $U_P(R) = (R-P)SF(P-R)$, where SF(x) is the unit step function, equal to 0 if x<0, and 1 thereafter. The generalized derivatives obey the following:

$$U'_P(x) = 1 - SF(x - P)$$
$$\frac{\partial}{\partial x} U'_P(x) = -\delta(x - P)$$
$$\frac{\partial}{\partial P}[U'_P(x)] = \delta(x - P) \quad ,$$

where in each case δ(·) is the Dirac delta function. Differentiating both sides of (11), we obtain

$$\text{(12)} \qquad \frac{\partial h}{\partial P} = \frac{E[f\delta(s - hf - P)]}{E[f^2\delta(s - hf - P)]} .$$

The numerator will be zero if and only if

$$\text{(13)} \qquad E[f\delta(s - hf - P)] = 0, \; a.s. \; in \; P.$$

Setting u = s-hf and u^0 = P yields condition (10ii). Condition (10i) follows from basic probability theory. Thus if p(f,u) denotes the joint density of f,u, then

$$E[f\delta(u - u^0)] = \int_{-\infty}^{\infty} fp(f, u^0)df = p_u(u^0) \int_{-\infty}^{\infty} fp(f|u^0)df .$$

Thus if (10ii) is true and $p_u(u^0) \neq 0$, then (10i) follows.

*Remark*s

(a) Condition (10) may be referred to as the expectational invariance of f with respect to u.

(b) The condition implies that E[fu]=0, so that the regressor and residual are uncorrelated. It will follow from this that the hedge ratio h is consistently estimable by means of OLS applied to the quasi regression (6). Thus if an optimal linear invariant hedge does exist, it will always be consistently estimable by applying OLS, though this may not necessarily be efficient as an estimator.

(c) Condition (10) does not imply *E[u | f] = 0*, i.e. conditioning in the other direction. Hence we cannot say that *E[s | f] = hf*, so that the condition does not amount to a statement about the conditional expectation. Nor is the condition equivalent to saying that regressor and residual have to be uncorrelated. One can have E[fu]=0 without condition (10) being satisfied, as some of the examples below will make clear.

(d) Expression (11) corresponds to the first order optimizing condition for $U_P(R)$. It follows from the spanning decomposition (9) that conditions (10i,ii) will also be sufficient for global invariance.

III Dependent volatility

It might be expected that linear hedges would be most efficacious when there exists an underlying linear structural model connecting s and f, though it should be remarked that condition (10) does not logically presuppose any such model. However, if we do suppose such a model, then we should look to the nature of its disturbance term for invariance of the linear hedge.

It is not necessary that model disturbances should be of the classical i.i.d form in order for invariance to hold. Consider, for example, the ARCH type specification,

$$s_t = \beta f_t + \omega(s_{t-1})\varepsilon_t,$$

where ε_t is specified as a zero mean random variable normalized to have unit variance, and independent of both f_t and s_{t-1} .It is also assumed that f_t and s_{t-1} are independent (e.g. the efficient market hypothesis). In this framework, we can ask whether the beta hedge will be invariant. Thus we can set $u_t = \omega(s_{t-1})\varepsilon_t$, effectively setting $h = \beta$. Given that ε_t and f_t are specified as independent,

$$E[f_t | u_t] = E[f_t | \varepsilon_t = \omega(s_{t-1})/u_t] = \int_{-\infty}^{\infty} f_t p(f_t) p_\varepsilon(\frac{\omega(s_{t-1})}{u_t}) df_t = p_\varepsilon(\frac{\omega(s_{t-1})}{u_t}) E[f_t] = 0, \ a.e. in\ u_t \quad ,$$

where $p_f(\cdot)$ and $p_\varepsilon(\cdot)$ indicate the respective marginal densities. Thus ARCH type effects do not invalidate the linear invariance of the econometric hedge.

Stochastic volatility

Things can change when the volatility depends upon contemporaneous variables, as in some stochastic volatility models. Contemporaneous volatility is a source of practical concern, for it means that the basis spread can change when the environment to be hedged becomes more stressful.

Condition (10) can be illustrated by assuming a structural relationship of the form

(14) $s = \beta f + \omega(f)\varepsilon,$

where ε is specified as a zero mean random variable normalized to have unit variance, and is independent of f. However, the volatility term ω(f) depends upon the current value of f (as everything is now contemporaneous, we shall simply drop time subscripts). Let $p_\varepsilon(\cdot)$ denote the density of ε and $p_f(\cdot)$ be the marginal density of f.. Also let h be a proposed linear hedge, so that u = s - hf = (β-h) f +ω(f)ε. Condition (10) can be expressed as

$$0 = E_{f,u}[f\delta(u-u^0)] = E_{f,\varepsilon}[f\delta(\varepsilon - \frac{u^0+(h-\beta)f}{\omega(f)})]$$
$$= \int_{-\infty}^{\infty} fp_\varepsilon(\frac{u^0+(h-\beta)f}{w(f)})p_f(f)df\ .$$

It follows that for the additive model (10), the hedge h = β will be invariant if and only if

(15) $E_f[fp_\varepsilon(\frac{u}{w(f)})] = 0$, almost everywhere in u.

A sufficient condition for (15) to hold is that the densities of f and ε are symmetric and that the volatility function ω(f) is also symmetric in f, i.e. ω(f) = ω(-f) . Thus symmetry favours invariance.

Asymmetric stochastic volatility

To see what can happen when condition (10) is violated, suppose that the volatility function in the additive model (14) is asymmetric. For instance, it could be that basis departures are more marked when returns s and/or f are substantial and negative, so that ω(f) is a diminishing function of f . One can readily see that this would be a source of concern. We can examine its effect by using the utility generator framework, and tracing the optimal hedge ratio h(P), as the generator node P varies along the real line.

If model (14) holds, the utility generator first order condition (11) can be written by using the iterated expectation, as

(16) $0 = E[fE_{\varepsilon|f}U_P'[(\beta-h)f + \omega(f)\varepsilon \le P]]$

$= E[f.prob[w(f)\varepsilon \le P - (\beta-h)f]]\ .$

Suppose that P=0 and the density of ε is symmetric about its mean of zero. Setting h=β will satisfy (16). Thus for model (14) the optimal hedge at P=0 is just h = β, which is effectively the econometric hedge. This gives us one point of reference in plotting the optimal hedge against values of P: we must have h(0) = β .

Indeed, one can say a bit more than this. From expression (12) above,

$$\frac{\partial h(P)}{\partial P} = \frac{E[f\delta(s-hf-P)]}{E[f^2\delta(s-hf-P)]}$$
$$= \frac{E[fp_\varepsilon(\frac{P-(h-\beta)f}{\omega(f)}]}{E[f^2 p_\varepsilon(\frac{P-(h-\beta)f}{\omega(f)}]}.$$

We observe that if the hedge ratio h is set equal to β and P=0, then the derivative of the OMD hedge function h(P) must be zero at P = 0. Thus the optimal hedge function must pass through the least squares hedge value at P = 0, and at that point must have a stationary value. Carrying out a further differentiation with respect to P, we find that the second derivative must be zero if the density p is symmetric. Hence the stationary point must be a point of inflexion.

Figure 2 plots the hedge function h(P) against P for the case where f is N(0,σ_f^2) and ε is N(0,1), with σ = 1.25% and β=1. It assumed that the volatility function is monotonically declining in f of the form $\omega(f) = w(0)2^{-f/\sigma_f}; w(0) = 0.4\%$. The hedge ratio is no longer invariant, but declines with P. However, the inflexion point at P =0 is also apparent. We can regard this as a flat area or ‘dead zone’ about the econometric hedge, and its implications are explored below.

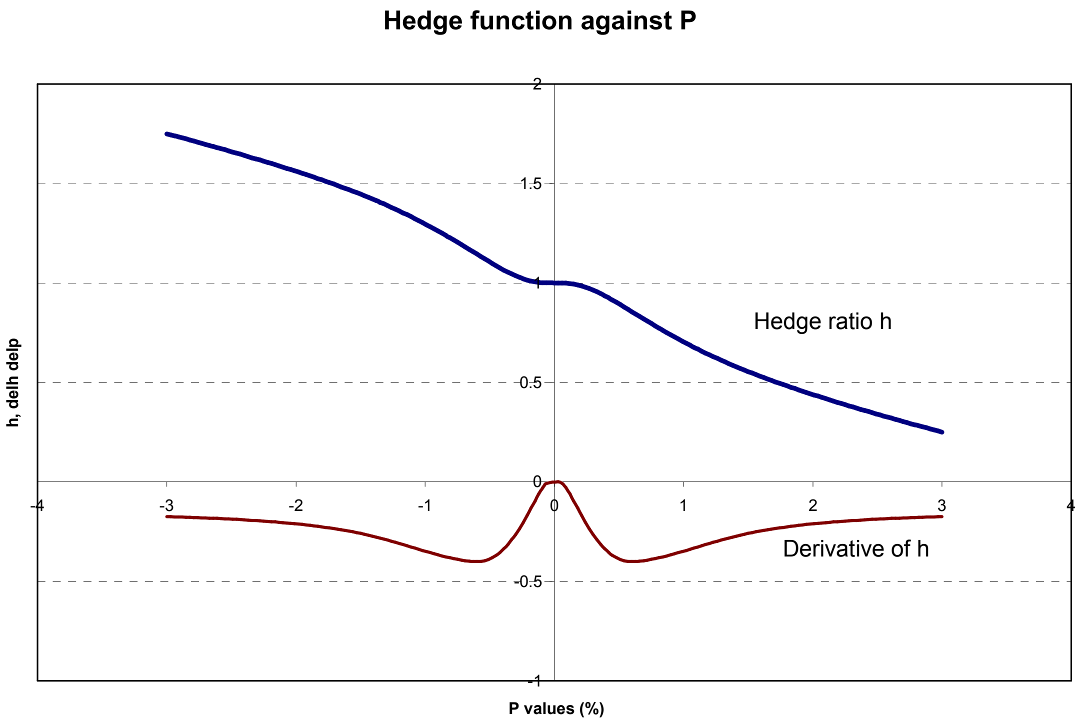

Figure 2 Point of inflexion to the hedge function

Hedge zone sensitivity

In general, the linear hedge will not be invariant. However, with OMD methods one can examine responsiveness to different areas along the return axis. For instance, the utility function depicted in figure 3 has a 'pain threshold' zone in the shaded area. The firm may be prepared to accept a return of zero with a degree of equanimity. But a loss of, say, 10% would be viewed as much more than proportionately significant. This might be characteristic of 'value at risk' decision environments (e.g. Wilson 1999), where the lower limit of capital availability is threatened. Above this zone, the curvature term (-U"(P)) in formulas (8) or (9) is small or negligible, and the same is true below the zone, by which time the damage is done. Hence most utility weighting will apply to the shaded zone. This means that the generator hedge solutions h(P) for $P \in Z$ will be those most relevant to the hedge decision.

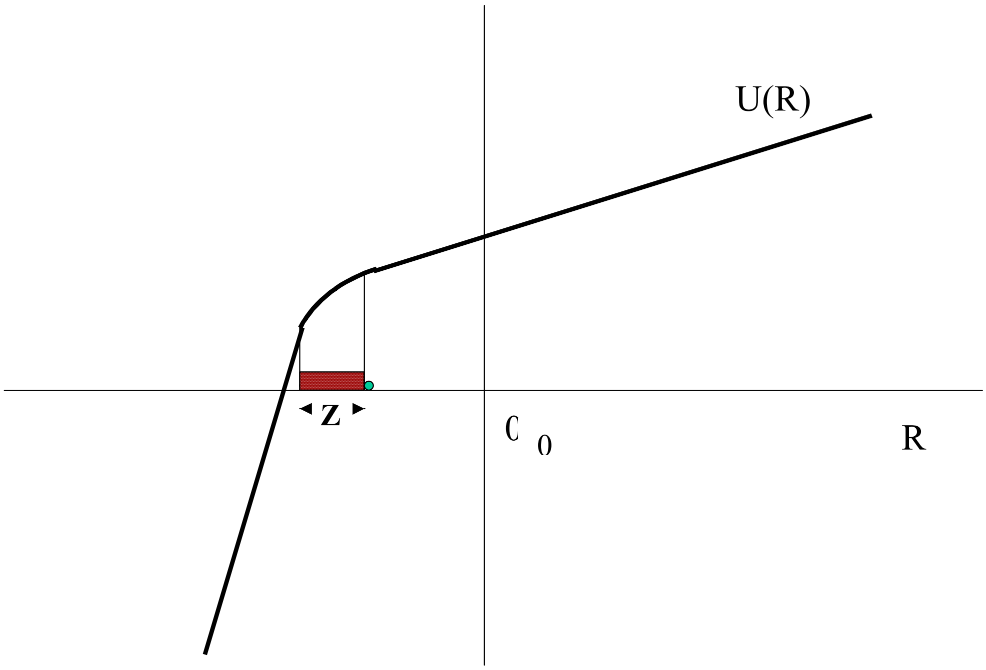

Figure 3 Local curvature and 'pain' thresholds

How could one advise a hedger on the basis of point of inflexion or 'dead zone' in the earlier figure 2? To do that one might call on the zones of distress argument implicit in the earlier figure 3. If the primary zone of distress is at or around zero return, then the advice would be to use the unadjusted econometric hedge, h = β. Even if the zone is a little beneath zero, the inflexion or 'dead zone' property of the econometric hedge means that little if any welfare would be lost. On the other hand if the zone of distress was heavily concentrated on low return values, say around P=-2%, then the recommendation would be to over hedge, adjusting h to above the value of the econometric hedge. The reason is that if basis volatility increases as f declines, low realisations of f are more likely to accompany even lower values of the natural position *s*.

IV Empirically based judgements

Although practical hedging decisions are theoretically concerned with conditional distributions, they have also to be grounded in empirical regularities. Thus in what follows, we suppose that the hedger has available a data window which he or she is reasonably confident will be stable, in the sense that the hedging relationship observed in the past will continue over the coming period. The preceding section suggests that the econometric beta hedge constitutes a reasonable starting point. Thus the task is to infer from the available data whether it is likely to be invariant and if not, whether it can feasibly be improved upon. The main difficulty in doing so is that one does not typically know the utility function, or at least have an exact enough specification for it for such a purpose. Thus is what follows we explore methodologies that are indicative in nature, perhaps involving some guided introspection as to preferences and their effects.

Directional methods

The OMD hedge schedule can be used as a directional indicator for the hedge ratio, relative to the econometric hedge. If $\tau(P)$ lies wholly above the horizontal axis, the interpretation is that increasing the hedge ratio would bring value to every risk averse agent. If the OMD schedule is positive only for low values of P, this would tell us that increasing the hedge ratio would have value for more risk averse agents, or for those in a pain zone that falls within the indicated P region.

Figures 4a-d illustrate. The data are daily return observations on the FTSE 100 index and associated LIFFE futures. It is divided into four equal sub periods as indicated , each of length 1135 trading days. One sigma confidence bands are attached as in Bowden (2000). The later periods all show evidence of positive OMD values around zero or less, indicating that sensitivity to negative outcomes (pain zones concentrated below zero) should raise the hedge ratio relative to the beta hedge. The first block exhibits contrary behaviour. The more risk averse agents should now moderate the hedge relative to the econometric beta hedge (in this case OLS). There is certainly no case for increasing it. At first sight this is a surprising conclusion as this block contains the Oct 19 crash of 1987. One might expect an increased demand for hedging to be revealed in the historical data. However figure 5 shows that the supposition would be premature. The vertical distance between the point and the line can be taken as a measure of the reward to hedging. This was certainly positive and substantial on Oct 19, but it was followed by a number of days (marked with arrows) on which the net rewards would have been large and negative. A flatter hedge line would have helped. This is as striking example of hedge dislocation ;

the hedger becomes exposed to adverse movements in the hedge basis (u), occurring in exceptional circumstances, which are just those times in which the hedge should have special value.

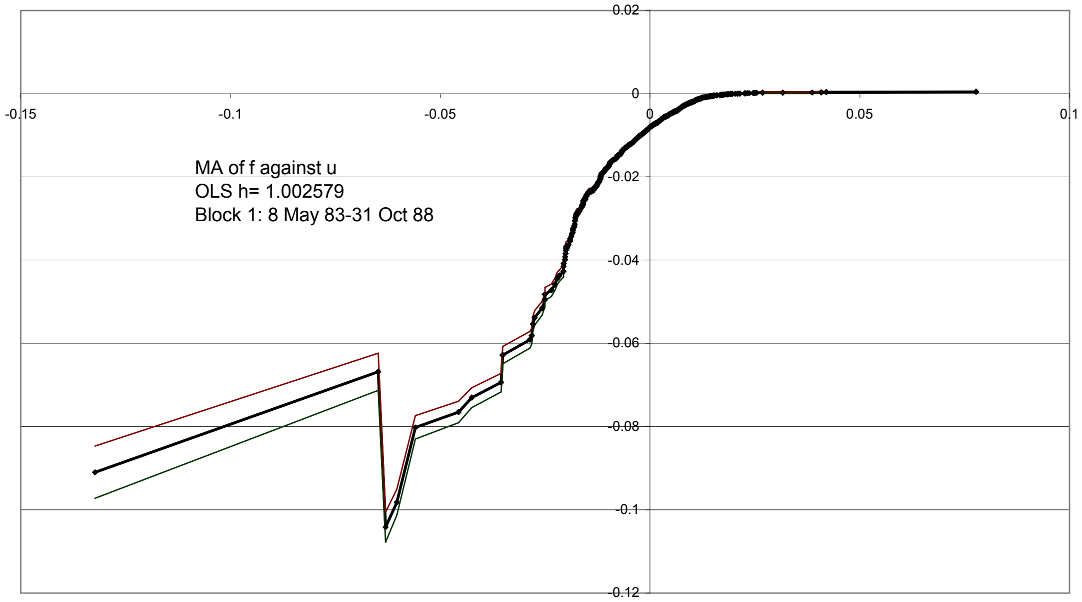

Figure 4a OMD plot, block 1

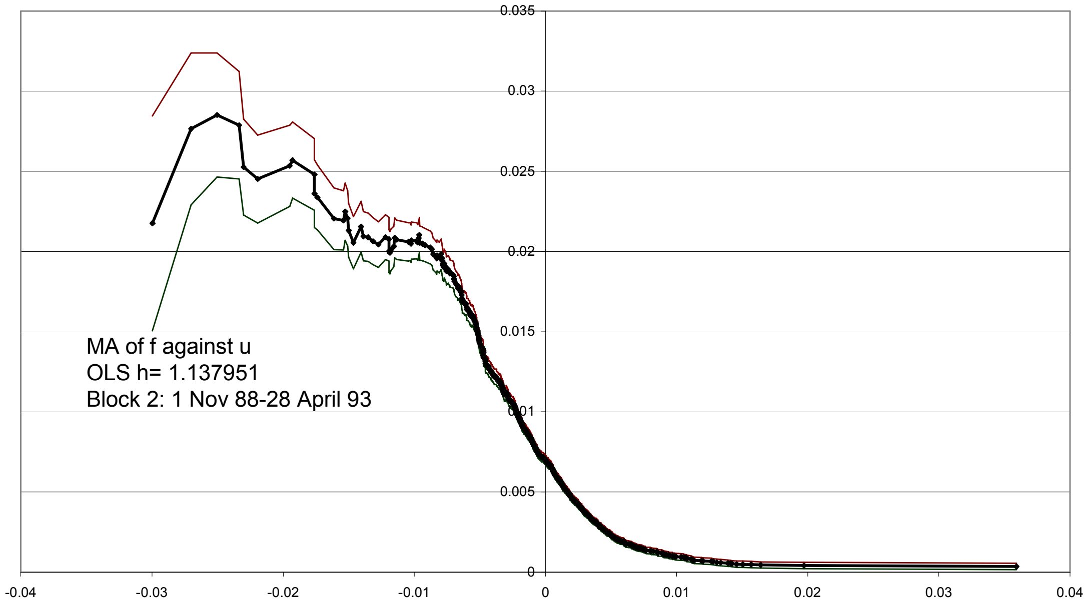

Figure 4b OMD plot, block2

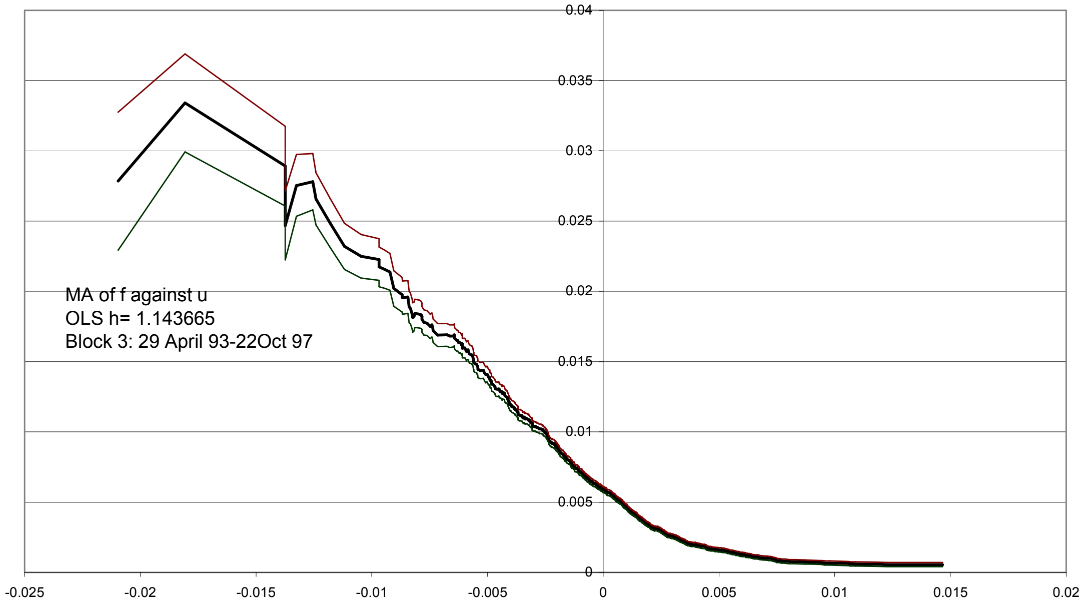

Figure 4c OMD plot block 3

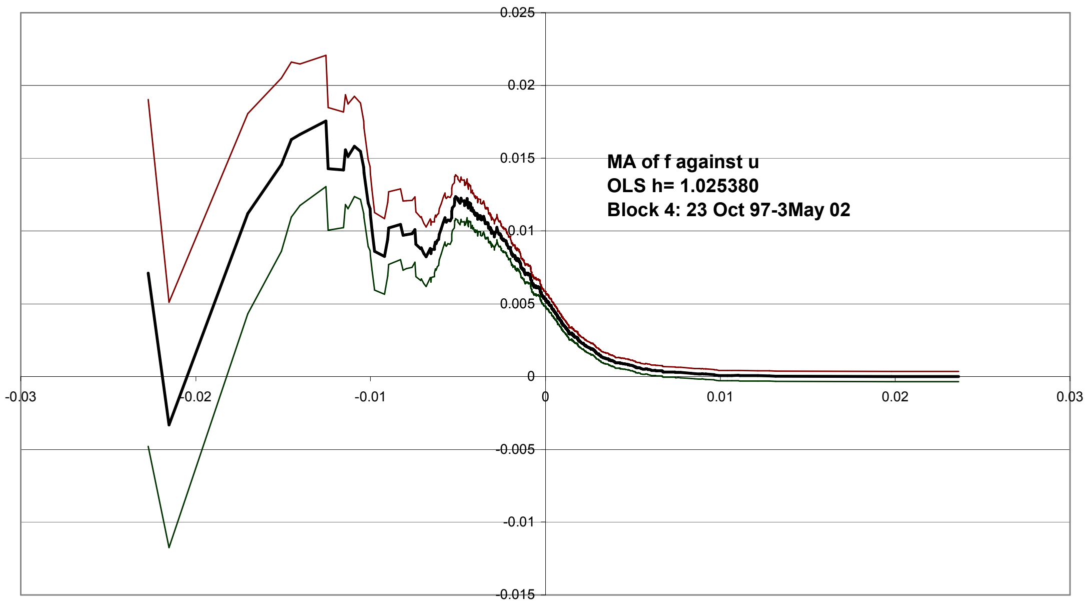

Figure 4d OMD plot, block 4

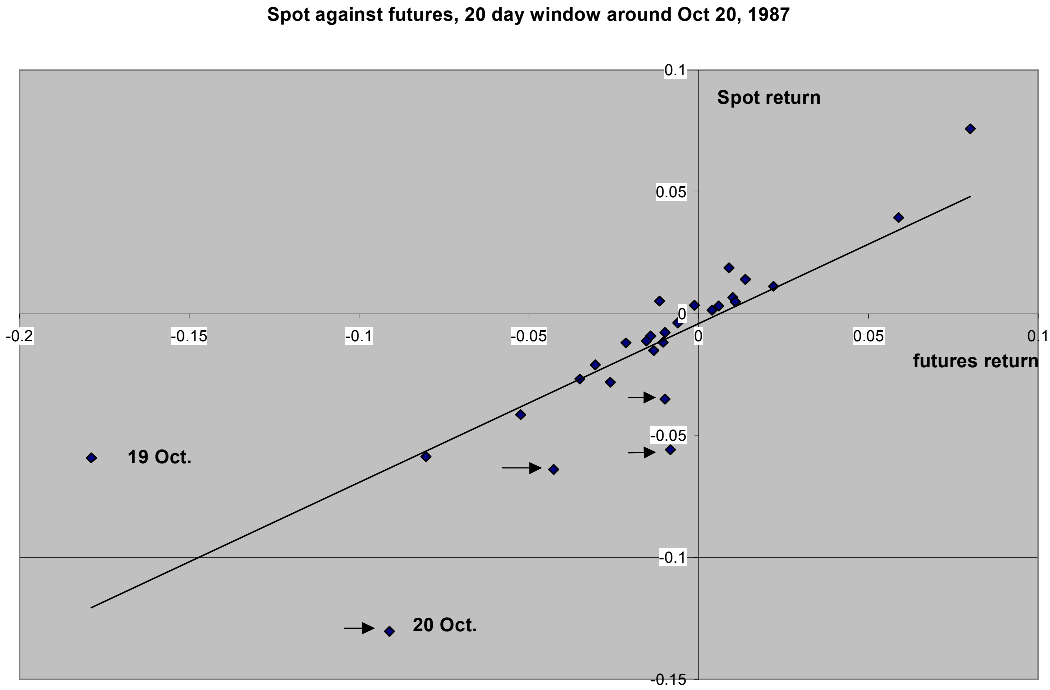

Figure 5 Returns around the 1987 crash

Plotting generator hedge ratios

An alternative is to plot the generator hedge ratios h(P) as a function of P. The optimand will be the sample expected value, namely $\frac{1}{T}\sum_{i=1}^{T} U_P(s_i - hf_i)$, and h is chosen to maximize this. The computation can be executed by means of a simple one dimensional search on the optimizing parameter h.

Operationally there are two difficulties to this procedure. First, if the optimand is rather flat, as it is under hedge invariance, then the routine will be uncertain just where to put h. Second, for low values of P there will not be too many observations generating non zero values for the optimand. The sample size factor [1/T] then ensures that their collective contribution will be extremely small and the effect of varying h will likewise turn out to be small. Hence this may not be a well conditioned estimation problem.

Figure 6 shows the optimizing values for the FSTE hedge data. For values of P at or close to zero there is a dead zone. Again we observe that block 1 differs from the others in that diminishing risk aversion perversely calls for a higher hedge ratio.

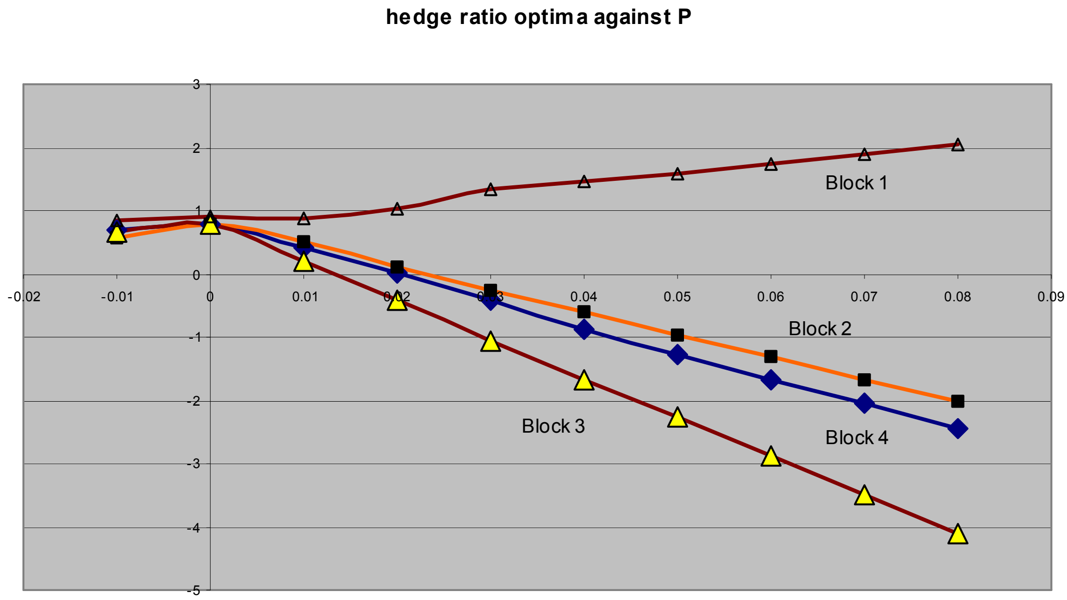

Figure 6 Hedge ratios plotted against P

Our general conclusion for the FTSE data is the optimal hedge ratio is not invariant. The linear econometric hedge offers insufficient protection to the manager who is sensitive to very low values of the spot return. However this has to be qualified by the nature of the underlying basis relationship. Given the last 3 periods, one would think that raising the econometric beta would raise welfare in the low return end of the axis but not overmuch diminish welfare at the upper end. However in the earlier period, which encompassed the exceptional 1987 crash, this would not have been a good policy. This is an instance where the regression relationship which forms the presumed starting point is not itself stable, and a proposed hedging rule has to have regard to the possibility that such episodes might arise, or loosely speaking, of ‘model uncertainty’.

IV Concluding remarks

Our substantive conclusions are the following:

(a) A condition is given for invariance, relating the hedge instrument return f to the basis residual u. Local invariance holds if and only if the hedge return is expectationally invariant with respect to the residual; this condition is sufficient also for global invariance. If this is the case then one can proceed to utilise the econometric beta, confident that it is the best for the purpose regardless of the precise nature of the hedging loss function.

(b) On the other hand, one cannot proceed on the basis that if classical regression assumptions hold then the econometric beta will necessarily be the best linear hedge. In other words, the converse to (a) is not necessarily true, and one can find important cases where regressor and error are uncorrelated but the hedge will not be invariant.

(c) The econometric beta is a good starting point, with the possible advantage of a zone of inflexion, wherein the optimal hedge is not much different.

(d) Ordered mean difference methods can be used to examine directional departures from the econometric beta. This can done in conjunction with judgement calls about sensitive zones of portfolio return. The OMD schedules of the hedge instrument against the provisional hedging basis can reveal whether the provisional hedge ratio should be raised or lowered.

(e) Judgment calls based on the above are usefully informed by an examination of the historical data for possible hedge dislocation, or basis risk.

References

Amemiya, T. (1982) 'Correction to a lemma', *Econometrica*, 50, 1324-28.

Antosik, P., J. Mikusinski and R. Sikorski (1973) *Theory of distributions: the sequential approach*, Amsterdam: Elseview/Warsaw:PWN.

Baillie, R. and R. DeGenaro, R. (1990) 'Stock returns and volatility', Journal of *Financial and Quantitative Analysis*, 25, 203-14.

Baillie, R. and R.T. Myers (1991) 'Bivariate GARCH estimation of the optimal hedge ratio for stock index futures: A note', *Journal of Futures Markets*, 15, 61-67.

Bollerslev, T., Engle, R.F. and J.M. Woolridge (1988) 'A capital asset pricing model with time-varying covariances', *Journal of Political Economy*, 96,116-327.

Bowden, R.J. (1995) 'Production organization and risk control when market instruments are available', *Management Science*, 6, 1073-82.

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

Bowden, R.J. (2002) 'Ordered mean difference benchmarking, utility generators and capital market equilibrium', School of Economics and Finance, Victoria University of Wellington, forthcoming *Journal of Business*.

Bowden, R.J. (2003) 'Utility spanning and the ordered mean difference envelope', revised, School of Economics and Finance, Victoria University of Wellington.

Brooks, C. and R. Chong (2001) 'The cross-currency hedging performance of implied versus statistical forecasting models', *Journal of Futures Markets*, 21,1043-69.

Gagnon, L. and G. Lypny (1995) 'Hedging short term interest rate risk under time varying distributions', *Journal of Futures Markets*, 15, 767-83.

Guermat, C. and R.D.F. Harris (2002) 'Robust conditional variance estimation and value at risk', *Journal of Risk*, 4, 25-41.

Harris, R.D.F. and J. Shen (2002) 'Robust estimation of the optimal hedge ratio', School of Business and Economics, University of Exeter.

Kroner, K.F. and J. Sultan (1991) 'Exchange rate volatility and time varying hedge ratios', *Pacific Basin Capital Markets Research,* 2,397-412.

Lighthill, M.J. (1959) *An Introduction to Fourier Analysis and Generalised Functions,* New York: Academic Press.

Lin, J.W., Najand, M., and K.Yung (1994) Hedging with currency futures: OLS versus GARCH, *Journal of Multinational Financial Management*, 4,45 67.

Mikosch, T. (2002) 'Modeling dependence and tails of financial time series', WP181 Laboratory of Actuarial Mathematics, University of Copenhagen, actuarial@math.ku.dk

Stein, C. (1973) 'Estimation of the mean of a multivariate normal distribution', in K. Kajek (ed) *Proceedings of the Prague Symposium on Asymptotic Statistics,* University Karlova, Prague.

Wilson, T. (1999) 'Value at risk', in C. Alexander (ed.) *Risk management and analysis*, Chichester, U.K.: Wiley.