Riskcenter.com
When Hedging Makes a Portfolio Worse
Eric Falkenstein
11/17/03

Risk management has often been called part art, part science. This cliché always bothered me, primarily because while I like to think I'm good at risk management, I know I'm not very artsy: I do not look good in black and most current art leaves me scratching my head. Hedging, however, is a nice topic to illustrate how risk management can be seen as a combination of art and science.

The first rule of hedging is 'do no harm.' Like no-fat cream cheese, there are some things that remove more good stuff than bad. A good hedge diminishes volatility proportionally *more than* the return, so that if the hedge reduces annualized volatility by X%, it should reduce annualized return by no more than X%. The best way to assess this is a simple Sharpe ratio test, does the hedge add, or subtract, to the Sharpe ratio?

A hedge can make a portfolio less attractive for two general reasons. First, it impacts the expected return too much. Good examples include taking the credit risk out of a convertible bond portfolio via asset swaps, a bank hedging out all of its interest rate risk, or a young man's retirement plan perfectly hedged with Spiders. As most hedges have rather straightforward expected returns, it should be obvious whether this is occurring. One key worth mentioning is that hedges should assume the long-run average return, not a conditional return—you do not want market timing bets embedded in your hedge positions (ie, hedging only when you think the market will decline is not hedging, it's market timing).

The second kind of bad hedge actually adds volatility to the portfolio. Clearly this would be unintentional or unexpected, but how could this happen? Primarily because a hedge worked in theory, but not practice, which means it had good backtests, but behaved differently in real time. This is the practical side of econometrics often not examined in the classroom, yet unstable parameters, the results of data-mining or regime changes, are the rule rather than the exception outside the ivory tower of academia.

The key to determining if a hedge will add to a portfolio's volatility in real-time is the t-stat in a univariate regression. If the t-stat is above 2, it should reduce volatility, if its below 2 it adds volatility and should be ignored. A t-stat is just the hedge ratio divided by its standard error, and so it is important to have the correct standard error. The most conspicuous standard error contaminant is caused by autocorrelation. Most time series are highly autocorrelated, and this vastly inflates their simple OLS standard errors. Most amateurs focus on the R^2, but the R^2 ignores degrees of freedom and only focuses on fit. I provide a proof below for the simple case with no autocorrelation.

I doubt that there have been any meaningful economic decisions that were primarily driven by a Chow test, a Wald test, or some manifestation of 2-stage least squares that was not apparent in ordinary least squares. But this statistic is different. In a *univariate* regression, a t-stat is highly informative and transparent, and therefore compelling. The art here is choosing the right criterion value for the t-stat, large enough to not constantly overwhelm common sense, while small enough to have practical significance (ie, the ability to change your actions). The actual t-stat criterion can vary depending as well on your Bayesian prior, and the degree to which the hedge seems theoretically appropriate, but at the very least it has you focusing on the key issue. If you want to take the hard nonBayesian counterposition, next time I'm in NYC we can meet at The Strand, smoke clove cigarettes, and discuss.

Proposition: Hedging lowers portfolio volatility when the t-stat is >2

Proof:

Let portfolio p have a relation to a hedge instrument h in the familiar OLS way

$$p = \beta h + \varepsilon \tag{1.1}$$

Let the hedge proxy of p be p*, h times its estimated hedge ratio β*, which can be thought of as the 'estimate' of the true portfolio p

$$p^* \equiv \beta^* h \tag{1.2}$$

The hedge ratio is estimated with error, specifically

$$\beta^* = \beta + \phi \tag{1.3}$$

$$s.e.(\phi) = s.e.\left(\beta^*\right) \sim N\left(0, \sigma^2_{\beta^*}\right) \tag{1.4}$$

the t-stat is defined as the ratio of the coefficient divided by its standard error

$$t - stat = \frac{\beta^*}{\sigma_\beta} \tag{1.5}$$

or

$$\sigma_{\beta^*} = \frac{\beta^*}{t} \tag{1.6}$$

Thus by eq (1.6), 2 times the standard error is therefore

$$2\sigma_\beta = \frac{2\beta^*}{t} \tag{1.7}$$

Now given that E(β*)=β, the measured hedge ratio is 'off' from the true ratio by this reasonably conservative two-standard deviation number

$$\beta^* = \beta \pm \frac{2\beta}{t} \tag{1.8}$$

The variance of the unhedged portfolio, p is

$$\begin{aligned}
&\operatorname{var}(p) = \operatorname{var}(\beta h \pm \varepsilon) \\
&\ldots = \operatorname{var}(\varepsilon) + \operatorname{var}(\beta h) \quad \text{(by assumption of OLS } h \text{ and } \varepsilon \text{ are independent)} \\
&\ldots = \operatorname{var}(\varepsilon) + \beta^2 \operatorname{var}(h) \\
&\ldots = \sigma_\varepsilon^2 + \beta^2 \sigma_h^2
\end{aligned} \tag{1.9}$$

The variance of the hedged portfolio, p-p*, is

$$\begin{aligned}
&\operatorname{var}(p - p^*) = \operatorname{var}\left(\beta h + \varepsilon - \left\{\beta + \frac{2\beta}{t}\right\} h\right) \\
&\ldots = \operatorname{var}\left(\varepsilon - \frac{2\beta h}{t}\right) \\
&\ldots = \operatorname{var}(\varepsilon) + \operatorname{var}\left(\frac{2\beta h}{t}\right) \\
&\ldots = \operatorname{var}(\varepsilon) + \frac{4\beta^2}{t^2} \operatorname{var}(h) \\
&\ldots = \sigma_\varepsilon^2 + \frac{4\beta^2}{t^2} \sigma_h^2
\end{aligned} \tag{1.10}$$

So the criterion of a hegde, that it at least reduce volatility, is that

$$\operatorname{var}(p - p^*) < \operatorname{var}(p) \tag{1.11}$$

using equations 1.8 and 1.9 we get

$$\sigma_\varepsilon^2 + \frac{4\beta^2}{t^2} \sigma_h^2 < \sigma_\varepsilon^2 + \sigma_h^2 \tag{1.12}$$

or

$$\therefore \quad t > 2 \tag{1.13}$$

QED