Appendix for "The Economic Value of Volatility Timing"

Suppose that our objective is to use the information contained in the return sequence $\{r_t\}_{t=1}^T$ to estimate the sample path of the unobserved volatility sequence $\{\sigma_t^2\}_{t=1}^T$. In general, we expect that the return realizations near t provide the most reliable information about the value of σ_t^2. Intuitively, therefore, we should focus on the data near t when estimating σ_t^2, and the accuracy of our estimates should increase as the time between successive observations of r_t shrinks. More formally, we can show that, under reasonable restrictions on the smoothness of the volatility process, a suitably-constructed rolling weighted average of the squared return innovations is a consistent and asymptotically normal estimator of σ_t^2.

Our rolling estimator of σ_t^2 is based on the weights that are asymptotically optimal in the absence of specific parametric information about the volatility process. It takes the form

$$\hat{\sigma}_t^2 = \sum_{l=-t+1}^{T-t} (\alpha_t/2) e^{-\alpha_t |l|} (r_{t+l} - \hat{\mu})^2, \tag{A1}$$

where $\hat{\mu}$ is our estimator of $\mathrm{E}[r_{t+l}]$ and α_t is the decay rate. Foster and Nelson (1996) show that the optimal choice of α_t is given by ϕ_t/θ_t, where ϕ_t^2 is the conditional variance of the increments to the σ_t^2 process and $(\theta_t^2 + \sigma_t^4)/\sigma_t^4$ is the conditional coefficient of kurtosis of r_t. Although this optimal decay rate generally time varies, we can eliminate the time dependency by making additional assumptions about the return generating process. Specifically, if we assume that the increments to σ_t^2 exhibit proportional volatility ($\phi_t = \phi\sigma_t^2$) and that the conditional coefficient of kurtosis is constant ($\theta_t = \theta\sigma_t^2$), then setting $\alpha_t = \phi/\theta$ is asymptotically optimal. The empirical evidence indicates that, at least for stock indexes, these are reasonable assumptions (see, e.g., French, Schwert, and Stambaugh (1987)).

Theory provides little guidance concerning an appropriate choice of ϕ. However, an important class of trading models implies that returns are conditionally normal (see, e.g., Clark (1973), Tauchen and Pitts (1983), and Andersen (1996)), which we can show implies $\theta = \sqrt{2}$. Thus, rather than simultaneously estimating both θ and ϕ, we assume that $\theta = \sqrt{2}$ and estimate ϕ by constructing the rolling difference sequence,

$$\delta_{t+n} = \frac{1}{n}\sum_{l=0}^{n-1}(r_{t+l} - \hat{\mu})^2 - \frac{1}{n}\sum_{l=1}^{n}(r_{t-l} - \hat{\mu})^2, \tag{A2}$$

which measures $(\hat{\sigma}_{t+n}^2 - \hat{\sigma}_t^2)$ for an appropriately defined, one-sided, flat weight rolling estimator. Given that $\phi_t = \phi\sigma_t^2$ and $\theta_t = \sqrt{2}\sigma_t^2$, it follows from Foster and Nelson (1996) that δ_{t+n} converges in distribution to a mean zero normal random variable with variance $(4/n + 2n\phi^2/3)\sigma_t^4$. Thus, by the law of large numbers,

$$\frac{3}{2n(T-2n)} \sum_{t=1+n}^{T-n} \left(\frac{\delta_{t+n}}{\hat{\sigma}_t^2}\right)^2 - 6/n^2 \to \phi^2 \tag{A3}$$

as the time between successive observations of the process goes to zero.

Now consider a rolling estimator of the conditional covariance matrix of a vector of returns. To develop comparable results, we must address some additional issues. There are no simple restrictions on a vector process that imply constant optimal decay rates for all elements of the covariance matrix. Moreover, the optimal decay rates generally differ across assets. This can cause problems when we try to invert the covariance matrix. We avoid these difficulties by developing a procedure that yields a single fixed exponential decay rate for all elements of the covariance matrix. Our approach is to estimate this parameter by minimizing the mean squared error (MSE) of the estimator of the variance of the conditionally efficient portfolio.

Assume, for the moment, that we know the portfolio weights that deliver the conditionally efficient portfolio at each point in time. In this case, we can find the optimal decay rate for our one-sided estimator of the covariance matrix. Specifically, we first compute the return on the conditionally efficient portfolio as $r_{pt} = \mathbf{w}_t' \boldsymbol{r}_{t+1}$ and use the approach outlined above to construct the optimal one-sided rolling estimator of its variance,

$$\hat{\sigma}_{pt}^2 = \sum_{l=-t+1}^{-1} \alpha_p e^{\alpha_p l} (r_{p,t+l} - \hat{\mu}_{p,t+l})^2. \tag{A4}$$

Then, we perform the minimization,

$$\min_{\alpha_{vc}} \sum_{t=2}^{T} (\hat{\sigma}_{pt}^2 - \mathbf{w}_t' \hat{\boldsymbol{\Sigma}}_t \mathbf{w}_t)^2 \tag{A5}$$

where

$$\hat{\boldsymbol{\Sigma}}_t = \sum_{l=-t+1}^{-1} \alpha_{vc} e^{\alpha_{vc} l} (\boldsymbol{r}_{t+l} - \hat{\boldsymbol{\mu}})(\boldsymbol{r}_{t+l} - \hat{\boldsymbol{\mu}})' \tag{A6}$$

denotes the one-sided rolling estimator of $\boldsymbol{\Sigma}_t$ based on the decay rate α_{vc}. Under our MSE criterion, this estimator is optimal in the class of rolling covariance matrix estimators that use a single fixed exponential decay rate.

In practice, the optimal portfolio weights are unknown, so we have to adopt an iterative procedure. We start by forming an equally-weighted portfolio of the assets and solving for the optimal value of α_{vc}. Once we have an initial estimate of the optimal α_{vc}, we use this estimate to construct a one-sided, exponentially-weighted, rolling estimator of the conditional covariance matrix of returns. This, in turn, allows us to form initial estimates of the weights that deliver the conditionally efficient portfolio at each point in time. We apply these weights, solve for a new estimate of the optimal α_{vc}, and iterate on the process until convergence is achieved.

References

Andersen, Torben G., 1996, Return volatility and trading volume: An information flow interpretation of stochastic volatility, *Journal of Finance* 51, 169–204.

Clark, Peter K., 1973, A subordinated stochastic process model with finite variance for speculative prices, *Econometrica* 41, 135–156.

Foster, Dean P., and Dan B. Nelson, 1996, Continuous record asymptotics for rolling sample variance estimators, *Econometrica* 64, 139–174.

French, Kenneth R., G. William Schwert, and Robert F. Stambaugh, 1987, Expected stock returns and volatility, *Journal of Financial Economics* 19, 3–30.

Tauchen, George E. and Mark Pitts, 1983, The price variability-volume relationship on speculative markets, *Econometrica* 51, 485–505.