On the Long Run Volatility of Stocks [1]

Carlos M. Carvalho
McCombs School of Business
The University of Texas

May 2015

[1] with Hedibert Lopes and Rob McCulloch.

Question

- What is the long-run variance of stock returns?

Stocks for the Long Run: Conventional Wisdom

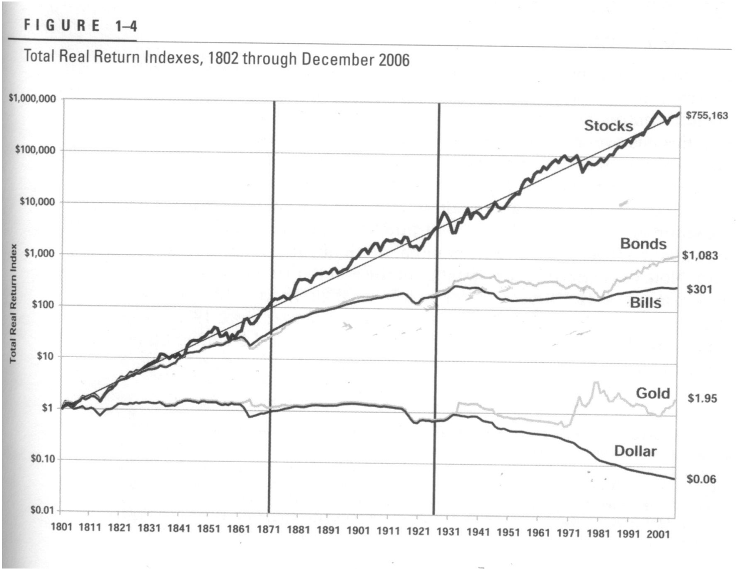

Stocks for the Long Run: Conventional Wisdom

- x-axis: Horizon y-axis: Volatility per year

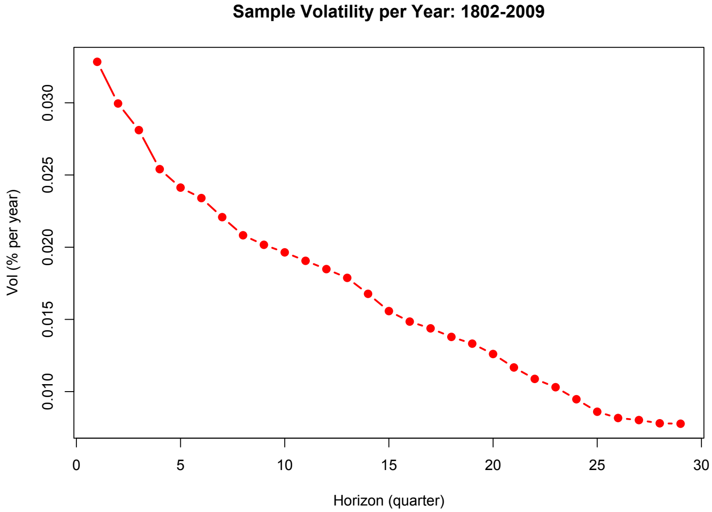

Stocks for the Long Run: Pastor and Stambaugh 2012 "main result"

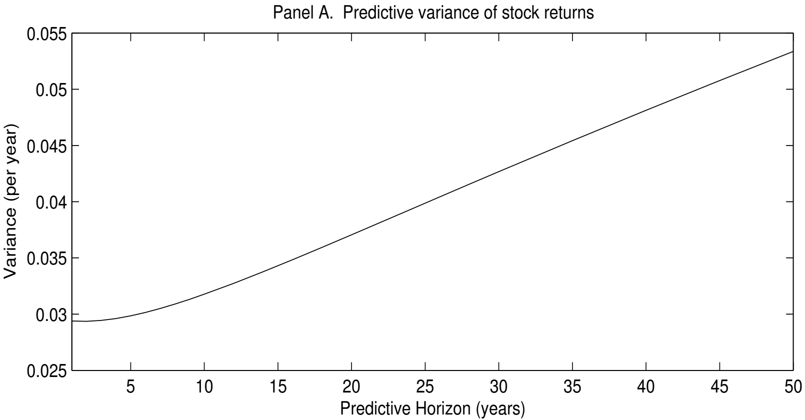

Summary

Taking a conditional approach, from the investor's perspective:

- A simple view of the world suggests that stocks are less volatile over long horizons (Barberis, 2000)...
- ...while a more complex view of the world states that stocks could be more volatile over long horizons (Pastor and Stambaugh, 2012)

Summary

Taking a conditional approach, from the investor's perspective:

- A simple view of the world suggests that stocks are less volatile over long horizons (Barberis, 2000)...
- ...while a more complex view of the world states that stocks could be more volatile over long horizons (Pastor and Stambaugh, 2012)

- Our work hopes to address:
  1. Which direction is right?
  2. Better understand the results sensitivity to prior specification
  3. Enrich PS2012 framework to include time-varying volatilities

Summary

Taking a conditional approach, from the investor's perspective:

- A simple view of the world suggests that stocks are less volatile over long horizons (Barberis, 2000)...
- ...while a more complex view of the world states that stocks could be more volatile over long horizons (Pastor and Stambaugh, 2012)

- Our work hopes to address:
  1. Which direction is right?
  2. Better understand the results sensitivity to prior specification
  3. Enrich PS2012 framework to include time-varying volatilities

- Our results indicates that I am not crazy for having 100% equity in my retirement portfolio, i.e., stocks are indeed less volatile in the long-run.

Background

- Returns k periods in the future:

$$r_{1,k} = r_1 + r_2 + r_3 + \cdots + r_k$$

Background

- Returns k periods in the future:

$$r_{1,k} = r_1 + r_2 + r_3 + \cdots + r_k$$

- If returns are i.i.d. $r_i \sim N(\mu, \sigma^2)$, i.e., $\approx$ *random walk* on prices:

$$Var(r_{1,k}) = k\sigma^2$$

  so that the variance per period is constant for any k investment horizon.

Background

- However, investor face parameter uncertainty...
- If $r_t \sim N(\mu, \sigma^2)$ and μ is unknown then,

$$\begin{aligned} Var(r_{t,t+k}|D_t) &= \mathrm{E}\left\{Var(r_{t,t+k}|\mu, D_t)\right\} + Var\left\{\mathrm{E}(r_{t,t+k}|\mu, D_t)\right\} \\ &= k\sigma^2 + k^2 Var(\mu|D_t) \end{aligned}$$

Background

- However, investor face parameter uncertainty...
- If $r_t \sim N(\mu, \sigma^2)$ and μ is unknown then,

$$
\begin{aligned}
Var(r_{t,t+k}|D_t) &= \mathrm{E}\{Var(r_{t,t+k}|\mu, D_t)\} + Var\{\mathrm{E}(r_{t,t+k}|\mu, D_t)\} \\
&= k\sigma^2 + k^2 Var(\mu|D_t)
\end{aligned}
$$

- Long run volatility (predictive variance) grows linearly with the horizon

Background

- If μ is mean reverting and

$$\begin{aligned} r_{t+1} &= \mu_t + u_{t+1} \\ \mu_{t+1} &= \alpha + \beta\mu_t + w_{t+1} \end{aligned}$$

where $Corr(u_{t+1}, w_{t+1}) < 0$,

$$Var(r_{t,t+k}) = k\sigma^2 \left[1 + 2A\rho_{uw} + B\right]$$

Background

- If μ is mean reverting and

$$\begin{aligned} r_{t+1} &= \mu_t + u_{t+1} \\ \mu_{t+1} &= \alpha + \beta\mu_t + w_{t+1} \end{aligned}$$

where $Corr(u_{t+1}, w_{t+1}) < 0$,

$$Var(r_{t,t+k}) = k\sigma^2 \left[1 + 2A\rho_{uw} + B\right]$$

- The effect of ρ_{uw} effect can dominate and imply a decreasing long run risk as both $A > 0$ and $B > 0$.

Background

- For stocks... using dividend yield as a proxy for expected returns:

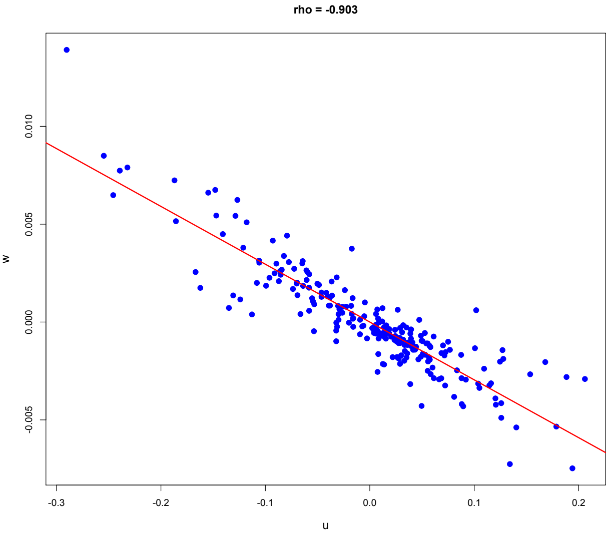

The General Model

$$
\begin{aligned}
r_{t+1} &= \mu_t + u_{t+1} \\
\mu_{t+1} &= \alpha + \beta\mu_t + w_{t+1} \\
x_{t+1} &= A + Bx_t + v_{t+1}
\end{aligned}
$$

where

$$
\begin{pmatrix} u_{t+1} \\ w_{t+1} \\ v_{t+1} \end{pmatrix} \sim N(0, \Sigma_{t+1})
$$

Priors and Posteriors... "weak prior" on ρ

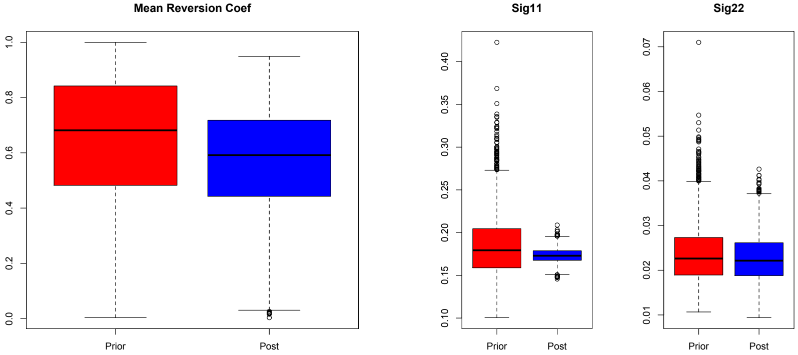

Priors and Posteriors... "strong prior" on ρ

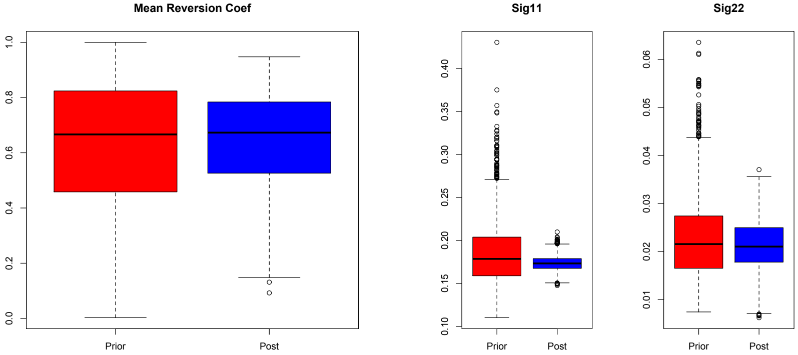

"Weak" vs. "Strong" Prior on ρ

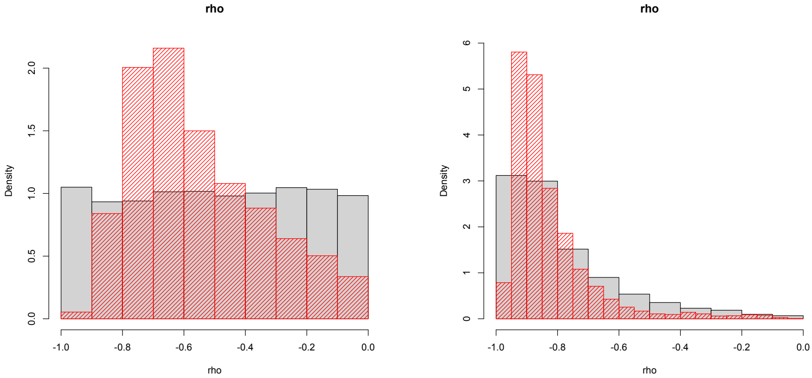

Figure: Histograms of draws from priors (gray) and posteriors (red) for ρ_{uw} in the "weak prior" (left) and "strong prior" (right) specifications.

Prior and Posterior Predictives

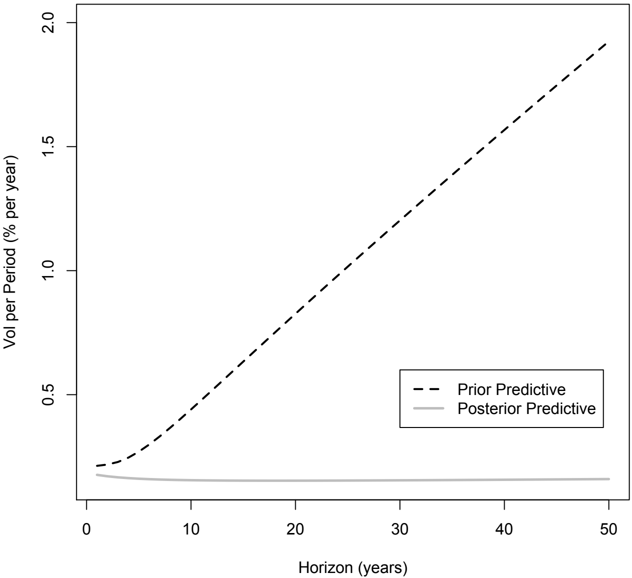

Long Run Volatilities per Period

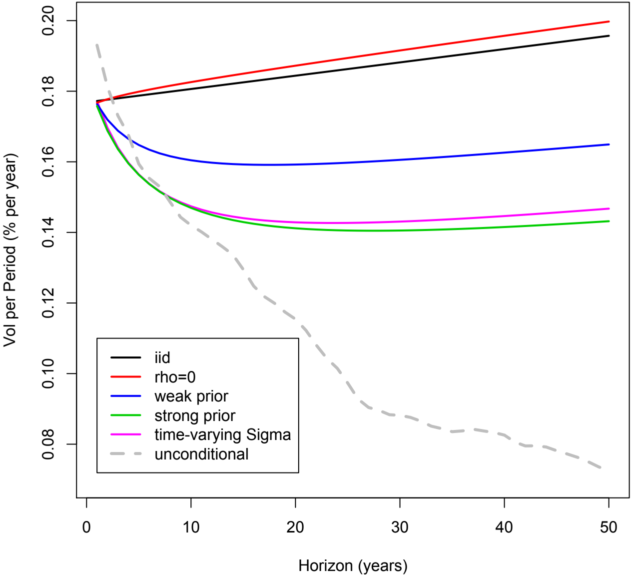

How Robust is the Result?

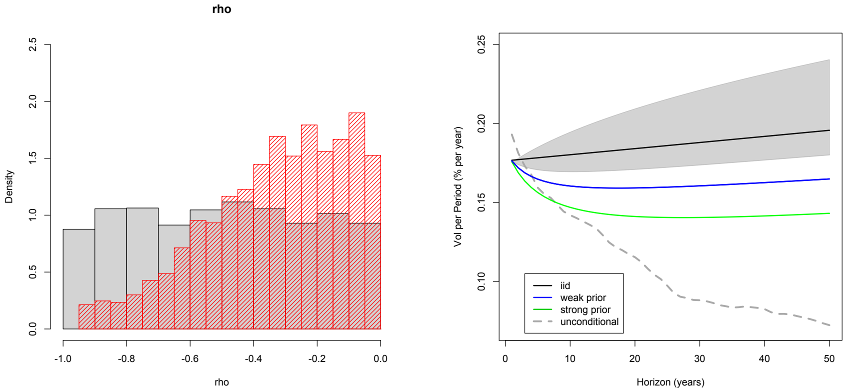

Figure: Histograms of draws from priors (gray) and posteriors (red) for ρ_{uw} in the "weak prior" set up where the 207 observations have been reshuffled so to break its dynamics.

Sources of Variance

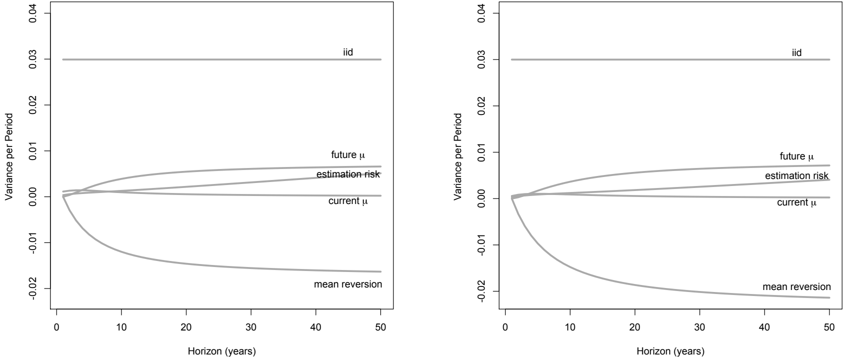

Figure: Decomposition of the predictive variance per period. The left panel is the results from the "weak prior" set up while the "strong prior" is in the right panel.

The Main Culprit!

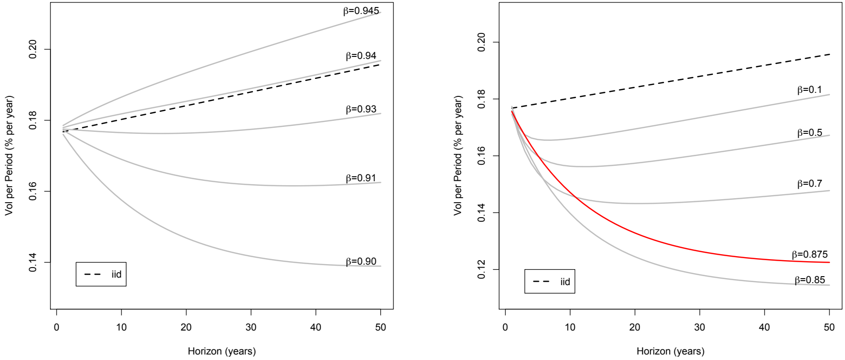

Figure: Predictive volatility per period plotted for different horizons for fixed values of β in the "weak prior" set up.

The Main Culprit!

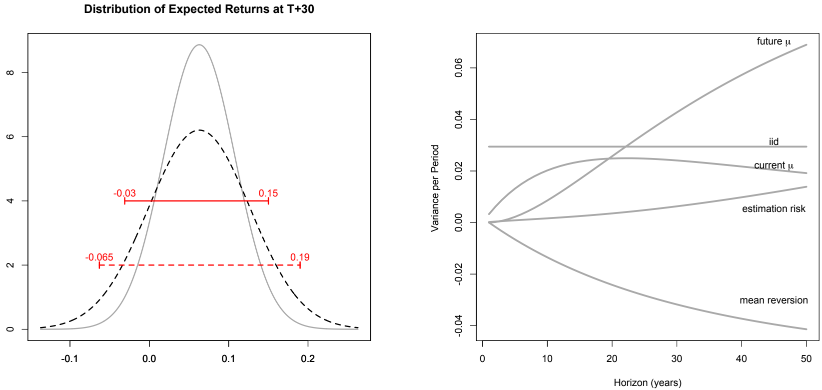

Figure: Left panel shows the posterior distribution of μ_{T+30} . The solid line in the right panel results from the "weak" prior set-up while the dashed line fixes $\beta = 0.945$. The right panel shows the decomposition of the predictive variance per period for the case where $\beta = 0.945$.

Portfolio Implications

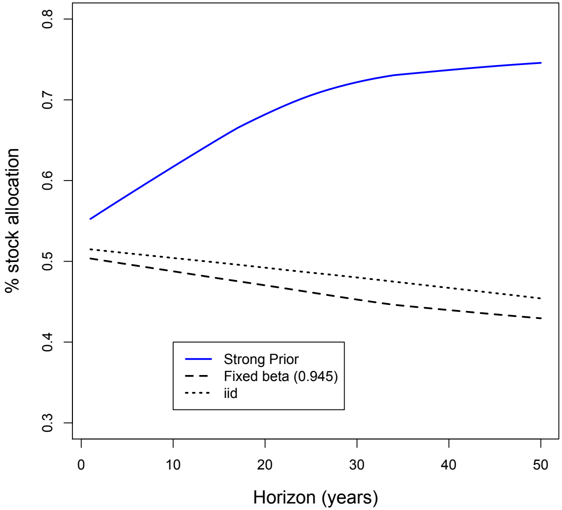

Replicating PS2012

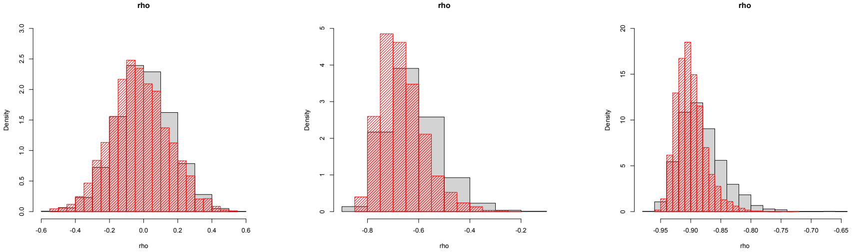

Figure: Priors (gray) and posteriors (red) draws of ρ_{uw} using priors from Pastor and Stambaugh 2012. In their terminology, from left to right: "non-informative", "less informative" and "more informative".

Replicating PS2012

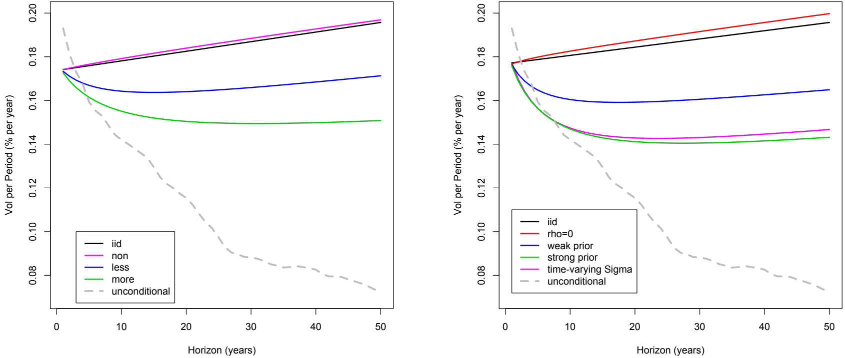

Figure: Comparison of our results (right) to the results using the priors in Pastor and Stambaugh 2012 (left).

Adding Predictors

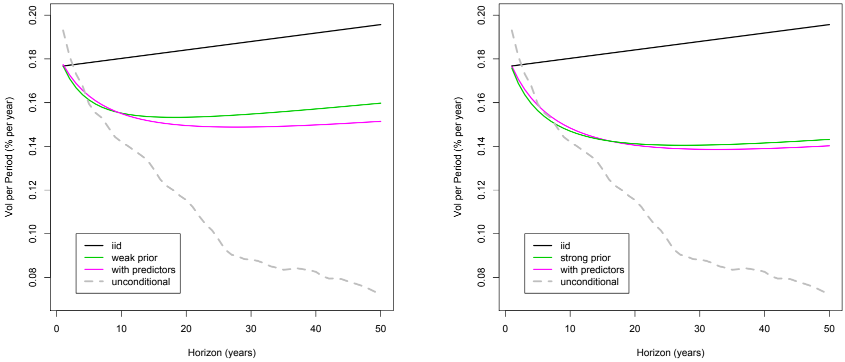

Figure: Predictive volatility per period plotted for different horizons when predictors are added. Results are for the "weak prior" (left) and "strong prior" set up.

Time-Varying Volatilities

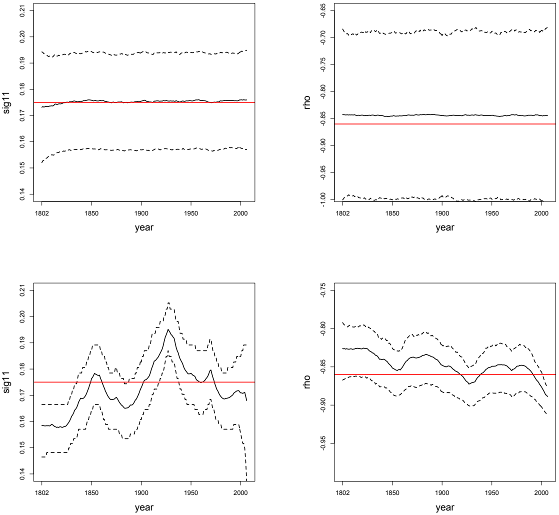

Closing Comments:

- With reasonable priors (or maybe even unreasonable) and 200 years of data stocks look very attractive for long horizon portfolios.
- This result appears robust to the added complexity of time varying relationships between predictors and expected returns and stochastic volatility.

Closing Comments:

- With reasonable priors (or maybe even unreasonable) and 200 years of data stocks look very attractive for long horizon portfolios.
- This result appears robust to the added complexity of time varying relationships between predictors and expected returns and stochastic volatility.
- The take home message is that conventional wisdom might not be so wrong after all...

Time Variation

Instead of just Σ, we want Σ_t, *and* we want to easily incorporate the prior belief that

$$\rho_t = corr(u_t, w_t) < 0, \text{for all } t$$

and possibly other prior beliefs as well.

Multivariate Stochastic Volatility

We start with the Choleski Stochastic Volatility approach of Lopes, McCulloch, and Tsay.

With one x we have:

$$\begin{array}{rcll} w_t & = & \exp(\theta_{t1}/2)\, Z_{t1} & p(w_t) \\ u_t & = & \theta_{t3}\, w_t + \exp(\theta_{t2}/2)\, Z_{t2} & p(u_t \mid w_t) \\ v_t & = & \phi_{t2}\, w_t + \phi_{t3}\, u_t + \exp(\phi_{t1}/2)\, Z_{t3} & p(v_t \mid w_t, u_t) \end{array}$$

At each t, the three θ's and three ϕ's are one to one with Σ_t.

Let's just focus on the θ's because they determine ρ_t.

Multivariate Stochastic Volatility

We have,

$$\begin{aligned} w_t &= \exp(\theta_{t1}/2)\, Z_{t1} \\ u_t &= \theta_{t3}\, w_t + \exp(\theta_{t2}/2)\, Z_{t2} \end{aligned}$$

$$\rho_t = \rho(\theta_{t1}, \theta_{t2}, \theta_{t3}) = \frac{\theta_{t3} \exp(\theta_{t1})}{\left[\theta_{t3}^2 \exp(\theta_{t1}) \times \exp(\theta_{t1})\right]^{1/2}}$$

Multivariate Stochastic Volatility

The usual prior for the θ_{ti} series is

$$\theta_{ti} = a_i + b_i \, \theta_{t-1,i} + s_i \, z_{ti}$$

Let's call this $q(\theta_{ti} \mid \theta_{t-1,i})$.

Letting $\theta_t = (\theta_{t1}, \theta_{t2}, \theta_{t3})$, let,

$$q(\theta_t \mid \theta_{t-1}) = \Pi_{i=1}^{3} q(\theta_{ti} \mid \theta_{t-1,i}).$$

We usually choose the s_i so that
successive θ are not "too different".

Prior Formulation

Our prior formulation is

$$p(\theta_t \,|\theta_{t-1}) \propto q(\theta_t \mid \theta_{t-1})\; f(\theta_t).$$

Prior Formulation

Our prior formulation is

$$p(\theta_t \,|\theta_{t-1}) \propto q(\theta_t \mid \theta_{t-1})\, f(\theta_t).$$

To get our ρ_t prior, we use,

$$f(\theta_t) = \exp\left\{ \frac{-(\rho(\theta_t) - \bar{\rho})^2}{\kappa} \right\}$$

Prior Formulation

Our prior formulation is

$$p(\theta_t \,|\theta_{t-1}) \propto q(\theta_t \,|\, \theta_{t-1})\; f(\theta_t).$$

To get our ρ_t prior, we use,

$$f(\theta_t) = \exp\left\{\frac{-(\rho(\theta_t) - \bar{\rho})^2}{\kappa}\right\}$$

q:
usual smoothness, don't let θ's jump around to much

f:
have preference for each θ_t, small κ means each θ_t should be such that $\rho_t \approx \bar{\rho}$

Bivariate Stochastic Volatility with Flexible Prior

$$(w_t, u_t)' \sim N(0, \Sigma(\theta_t)), \quad \theta_t = (\theta_{t1}, \theta_{t2}, \theta_{t3})$$

$$\begin{aligned} w_t &= \exp(\theta_{t1}/2)\, Z_{t1} \\ u_t &= \theta_{t3}\, w_t + \exp(\theta_{t2}/2)\, Z_{t2} \end{aligned}$$

$$\begin{aligned} p(\theta_t \,|\theta_{t-1}) &\propto q(\theta_t \,|\, \theta_{t-1})\; f(\theta_t) \\ &= q(\theta_t \,|\, \theta_{t-1})\; f(\theta_t)\; K(\theta_{t-1}) \end{aligned}$$

$$p(\theta_0) \propto f(\theta_0)\; \Pi_{i=1}^{3} p(\theta_{i0})$$

Simple Example

Let w and u be the observed bivariate series consisting of daily returns from two stocks in the S&P100.

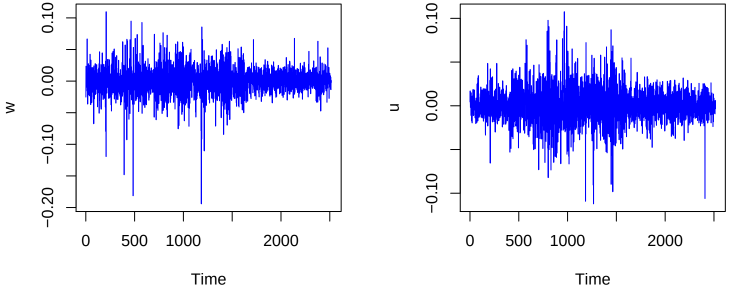

Prior:

$$f(\theta_t) = \exp\left[\frac{-\left(\rho(\theta_t) - \bar{\rho}\right)^2}{\kappa}\right]$$

For this data, it is more reasonable to believe that $\rho_t > 0$!

I'll hide the details about q and show results for

$$\bar{\rho} = .8, \;\; \kappa = .01, \; .25$$

$\kappa = .01$: tight prior.

$\kappa = .25$: loose prior.

loose prior: draws from prior

black is average draw, others are individual draws
(1,1): ρ_t, (1,2): θ_{t1}
(2,1): θ_{t2}, (2,2): θ_{t3}

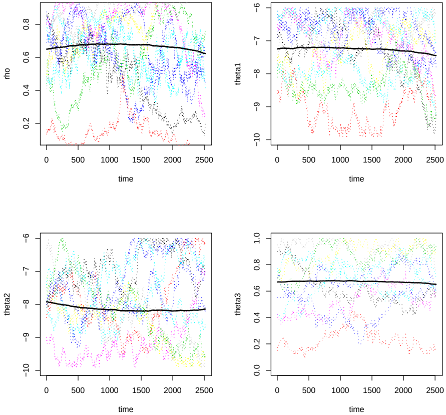

loose prior: draws from posterior

black is average draw, others are individual draws
(1,1): ρ_t, (1,2): θ_{t1}
(2,1): θ_{t2}, (2,2): θ_{t3}

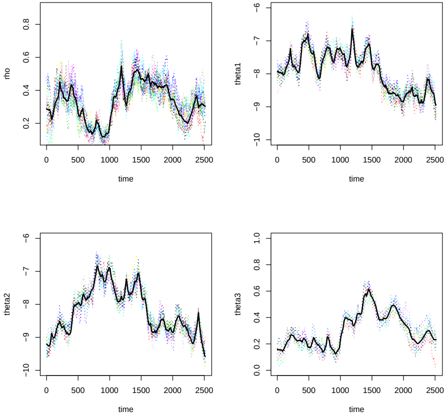

tight prior: draws from prior

black is average draw, others are individual draws
(1,1): ρ_t, (1,2): θ_{t1}
(2,1): θ_{t2}, (2,2): θ_{t3}

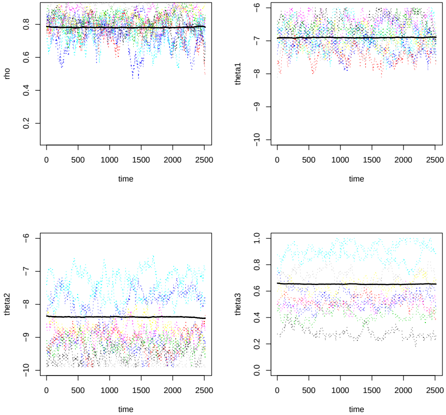

tight prior: draws from posterior

black is average draw, others are individual draws
(1,1): ρ_t, (1,2): θ_{t1}
(2,1): θ_{t2}, (2,2): θ_{t3}