A Bayesian View of Temporary Components in Asset Prices

Bjørn Eraker*

This draft: February 2007

Abstract

This paper studies models in which the a stock price contains a random walk and a stationary component, as in Fama & French (1988) and Poterba and Summers (1988). We extend this model to allow for two latent factors which generate short term and long term autocorrelations, respectively. To facilitate econometric identification, we assume that these factors are common across multiple asset returns, and we estimate the factor loadings. In an application to size and book/market sorted portfolios, we find the short term factor economically and statistically insignificant. Estimates of parameters relating to the long range component suggest that portfolios of small firm stock display about three times the amount of mean reversion than for large firm stocks. Overall, the evidence suggests that mean reversion is largely a small firm phenomenon. The evidence is consistent with dynamic equilibrium models in which asset prices co-integrate with aggregate consumption or dividends.

*bjorn.eraker@duke.edu. I am grateful for comments from Michael Brandt, Tim Bollerslev, George Tauchen and seminar participants at the Duke Financial Econometrics workshop. A previous version of the manuscript was circulated under the title "Predicting Asset Returns using Autoregressive Latent Component Models."

1 Introduction

Few topics in finance have been debated as heavily as the random walk hypothesis of security prices. There is no doubt that a large fraction of the daily or annual stock market return is unpredictable, as the random walk hypothesis would suggest. While it is understood that asset returns contain a very small predictable component at best, interest in the issue persists because even small amounts of predictability could prove valuable to asset managers.

In the debate over predictability, it is widely recognized that there exists time series such as price-multiples, interest rates and other macro variables that correlate significantly with future asset returns using standard methods of statistical inference. The results are still controversial because ”standard methods of inference” do not adjust for the fact that these variables are selected out of a larger set of variables which leads to an overstated level of statistical significance (”data-snooping bias”). A second strand of the literature focusses on serial correlations. Results presented in Fama and French (1988) in support of the predictability hypothesis, showing seemingly significant long run autocorrelations consistent with a model in which stock prices contain a random walk, and a stationary component. This evidence is widely contested. Richardson (1989) shows that the Fama-French analysis understate the standard errors in long term autocorrelations. Also, the infamous U shaped pattern in the autocorrelation function disappears once the highly volatile 1926 to 1940 period is removed from the data set. Fama (1991) concludes that the tests based upon the autocorrelation function have low power due to the effectively small sample sizes for long horizon correlations.

The lack of power in autocorrelation based tests poses a challenge which cannot be resolved unless the econometrician is willing to impose parametric assumptions about the functional form of the autocorrelation function. This can be accomplished by assuming a particular form of the dynamic model that generates asset returns. For example, in a model in which prices and dividends are cointegrated, as is the case in many dynamic equilibrium models, stock prices follow the stationary component model studied by Fama French (1988) and Poterba and Summers (1998). The model (henceforth the Shiller-Summers model as coined in Fama (1991)) has only four parameters which jointly describe the autocorrelation function.

A generalized version of the Shiller-Summers model is estimated in Lamoureux and Zhou (1996). Their paper studies a model in which the stationary price component follows an AR(4) process. In a Bayesian analysis of the structural model, Lamoureux and Zhou conclude that the data contain little or no evidence of a predictable component. Their analysis also reveal that the data are very uninformative about the structural parameters. Many of the autoregressive coefficients that drive the stationary price component have posterior standard deviations which are almost indistinguishable from the prior standard deviations.

In this paper we study a generalized version of the Shiller-Summers model. We generalize the model in two important ways. First, the price process has three components: A permanent (random walk) shock, a long run mean reverting component and a short term predictable component. Second, the three factor structure makes the model very hard to identify from univariate time series data. To overcome the identification problem, we propose a model structure in which asset returns are generated in a linear factor form: the returns on all assets are linear functions of a short term AR(1) process and a long range differenced AR(1) process, as well as asset specific shocks. The coefficients relating the individual asset returns to the common factors are estimated jointly with the parameters governing the dynamics of the state variables. The class of models considered here are perfectly consistent with dynamic equilibrium models in which the time-variation in which the price responds negatively to a temporary increase in risks.

The hypothesis that prices contain a permanent, random walk component and a mean reverting component is perfectly consistent with equilibrium models. In particular numerous models produce equilibrium price/dividend dynamics of the form $P_t/D_t = G(F_t)$ where G is some nonlinear function of a set of state variables, F_t. The state variables are typically related to time varying risks. For example, a model with time-varying volatility of the intertemporal marginal rate of substitution will generate this type of dynamics with F_t equal to the volatility factor. Under additional assumptions, $G()$ can be made to be exponential affine or approximately exponential affine. This is the case in Campbell and Shiller (1987,1988), Bansal and Yaron (1997), Bansal, Dittmar and Lundblad (2003,2004), among others. For these models, the equilibrium log price equals the log dividend plus the stationary component F_t and thus produce equilibrium price dynamics which is of the form considered in this paper. From this point of view, believers in efficient markets should not be surprised to find that stock prices contain some long term predictable component. In fact, random walk stock prices is not consistent with the dynamic equilibrium models in the above referenced papers.

A summary of the results obtained in this paper are as follows. First, we fit the Shiller-Summers model to size and book sorted portfolios using univariate specification. The parameters are quite well identified using a uniform (uninformative) prior distribution. Using the univariate model specification, we find moderate amounts of predictability. The predictable component is not deemed statistically significant. The lack of precision in parameter estimates, and thus also the latent extracted factor, lead to a significant posterior probability mass on the possibility that returns contain essentially only a random walk component. In moving to our multivariate generalization of the Shiller-Summers model, we obtain much sharper estimates of the parameters of interest. The results show some interesting patterns. Predictability is statistically significant at the three to ten year horizon for portfolios of small firms. Large firm portfolios, on the other hand, all contain insignificant amounts of predictability. The "factor loadings" increase almost uniformly along the size dimension.

Our multivariate model specification also contain a common factor designed to capture predictable variation in returns at short horizons. Specifically, the "short term factor" is an AR(1) process that enter linearly in the return generating process. This is consistent with the equilibrium specifications in Bekaert and Harvey (1995), Brandt and Kang (2004), and Johannes, Polson and Stroud (2004), among others. Our analysis does not find that this short term factor generates significant predictability in returns: The autoregressive coefficient that determine the amount of predictability has most of its mass near zero. Interestingly, the factor loadings associated with this short term component are widely dispersed along the book/market dimension consistent. This is consistent with permanent differences in risks across book/market assets.

The remainder of the paper is organized as follows. Section two discusses the model setup, introduces a generalization of the basic two component model and discusses estimation. Section three presents the empirical evidence and section four concludes.

2 Model

2.1 Basic model

Let p_t be the log value of a portfolio in which dividend payments are re-invested such that $r_t = p_t - p_{t-1}$ is the log-return. The model studied in Fama and French (1988), Poterba and Summers (1988) is

$$p_t = q_t + z_t, \quad (1)$$

$$q_t = q_{t-1} + \mu + \sigma_\eta \eta_t, \quad (2)$$

$$z_t = \phi z_{t-1} + \sigma_\epsilon \epsilon_t. \quad (3)$$

Likelihood inference for this model is possible under distributional assumptions $\eta_t, \epsilon_t \sim iid.N(0,1)$. The properties of the model are well documented, but we repeat some basic features here for completeness. The long run expected rate of return is μ. The process z_t is unobserved by the econometrician. The expected T period return at time t conditional upon the value of z_t, is $E_t(p_{t+T} - p_t) = E_t(z_{t+T} - z_t) + \mu T = (\phi^T - 1)z_t + \mu T$.

Thus, we have that the expected rate of return is less than the unconditional expected return μ whenever $z_t > 0$, and vise versa. In other words, when the value of the unobserved component z_t exceeds its average long term value of zero, the expected rate of return drops below its long term mean μ. The speed of mean reversion in the stationary component, ϕ, determines the shape of the autocorrelation function. For small values of ϕ, the model displays short term negative autocorrelation, while values of ϕ close to unity lead to long term reversals. A unit value of ϕ implies that the returns are random walk[1].

The properties of the Shiller-Summers model are well documented. Fama and French (1988) derive the autocorrelation function. In the following, we provide some additional insights into the dynamic behavior of the model. Define

$$R^2(T) = \frac{\mathrm{Var}(E_t(p_{t+T} - p_t))}{\mathrm{Var}(p_{t+T} - p_t)}$$

to be the fraction of variance of the expected rate of return to the total variance. The definition follows the standard definition of an R-square from linear regression. We can show that

$$R^2(T) = \frac{(\phi^T - 1)^2}{\phi^{4T} - 3\phi^{2T} + 2 + \frac{\sigma_\eta^2}{\sigma_\epsilon^2} T(1 - \phi^2)}. \tag{4}$$

Figure 1 illustrates how $R^2(T)$ changes with different values of the key parameters ϕ and σ_ϵ. The left plot keeps $\sigma_\epsilon = \sigma_\eta = 0.2$ (fixed) and illustrates that the R^2 is increasing and/or decreasing over various values of ϕ. A value of ϕ close to zero always implies more predictability at the short horizon. The R^2 approaches zero as ϕ approaches unity - the case in which returns are independent. Note that the R^2's may be large at longer forecasting horizons even when ϕ is fairly close to unity. This is true for relatively large values of σ_ϵ. The right plot in figure 1 illustrates that the R^2

[1] In this case, we cannot separately identify σ_η and σ_ϵ. That is, the likelihood function has a ridge at $\sigma_\eta^2 + \sigma_\epsilon^2 = \hat{\sigma}^3$ where the latter is the ML estimator for σ^2 under the assumption $r_t \sim N(\mu, \sigma^2)$. Notice that even though Maximum likelihood estimators for the two variance components are not well defined when $\phi = 1$, Bayesian posterior distributions are well defined because the posterior distribution of $\sigma_\eta^2 + \sigma_\epsilon^2$ is proper on $\mathbb{R}_+^2$.

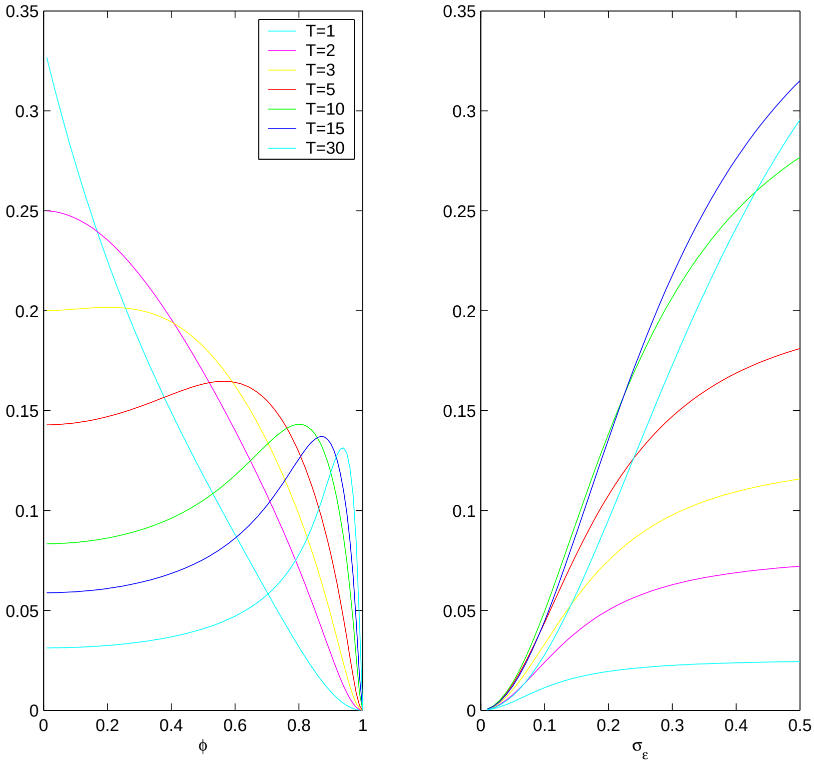

Figure 1: The figure shows $R^2(T)$'s as functions of ϕ (left) and σ_ϵ for various forecast horizons.

approaches zero as σ_ϵ approaches zero. Thus, even with the prior on $p(\phi) \sim U(-1, 1)$ which necessarily leads to an estimate of ϕ strictly less than unity, a low value of σ_ϵ will imply little or no predictability.

2.2 Multivariate generalization

The development of a multivariate generalization of the Campbell-Shiller model is motivated by the lack of conclusive empirical evidence from the univariate model. Estimates of the posterior distributions of the relevant model parameters ϕ, σ_ϵ and σ_η reported in the next section are sufficiently dispersed that when mapped into the corresponding posterior for R^2, it becomes clear that the posteriors for the R^2's place imply a substantial probability of no predictability.

In the following we consider a model that maintains the basic structure of the Campbell-Shiller model, but is significantly more parsimonious. This model is

$$\begin{aligned} r_t &= \beta \triangle z_t + \gamma f_t + \mu + \eta_t, & \eta_t &\sim N(0, \Omega), & (5) \\ z_t &= \phi z_{t-1} + \sigma_\epsilon \epsilon_t, & \epsilon_t &\sim N(0, 1), & (6) \\ f_t &= \kappa f_{t-1} + \sigma_f e_t, & e_t &\sim N(0, 1), & (7) \end{aligned}$$

where r_t is a vector of N returns. β, γ are length N vectors containing "factor loadings" associated with the latent factors z_t and f_t.

The covariance matrix Ω of asset specific shocks to returns is modelled as

$$\Omega = DRD$$

where D is matrix with diagonal element i, i given by σ_i, and hence, D contains asset specific standard deviations. R is a correlation matrix. Modelling the correlation matrix R constitutes a significant challenge because the number of parameters needed to model pairwise, unrestricted correlations grow quadratically in the number of assets. For example, for the twenty five Fama-French five by five size/book sorted portfolios we would need three hundred correlation coefficients. This would likely be an unfeasible modelling task. Lack of parsimony in the modelling of Ω would aversely affect the identification of the structural parameters. For this reason, we simply assume that the correlation between all shocks to the assets is the same, determined by a parameter ρ which is to be estimated.

The major difference in the multivariate model above, and the univariate Campbell-Shiller model is that ours has two single latent variables whereas in the C-S model, there is one latent variable per observed time-series. Our model still maintains the basic feature of the C-S model in that each asset price series has a random walk and a stationary component.

We can derive R^2, as defined above. Some algebra produces

$$
\begin{aligned}
R^2(T) &= \frac{\gamma^2 V_e + \beta^2 Q_e}{\gamma^2 V_T + \beta^2 Q_T + T\sigma_\eta^2} && (8)\\
Q_e &= \frac{(\phi^{2T}-1)^2}{(1-\phi^2)}\sigma_\epsilon^2 && (9)\\
Q_T &= (\phi^{4T} - 3\phi^{2T} + 2)\frac{\sigma_\epsilon^2}{1-\phi^2} && (10)\\
V_e &= \frac{\sigma_f^2 \kappa^2 (1-\kappa^{2T})}{(1-\kappa^2)^2} && (11)\\
V_T &= \frac{\sigma_f^2(\kappa^2 T - 2\kappa^{T+2} + \kappa^{2(1+T)} + \kappa^2 - 2\kappa^{T+1} + 2\kappa - T)}{(\kappa+1)(\kappa-1)^3}. && (12)
\end{aligned}
$$

The expression for the R^2 consists for four parts where Q_e and Q_T measure the expected and total variation in the z_t component, respectively. Similarly, V_e and V_T measure the expected, total variation in f_t. The asset specific parameters γ and β measure the sensitivity of a particular asset to the two predictable components. When these coefficients are zero, there is no predictability and $R^(T) = 0$ for all T.

2.3 Estimation

Estimation of the models considered here is straightforward. Indeed, the simple linear form leads to an easy adaptation of Kalman filtering techniques, as discussed below. This is one of the main advantages of the model setup here, as opposed to, say, a generalized model which allows for stochastic volatility. Stochastic volatility models require numerical integration techniques which are computationally costly.

A standard multivariate unobserved components model take on the form

$$
\begin{aligned}
y_t &= h' x_t + e_t && (13)\\
x_t &= F x_{t-1} + u_t && (14)
\end{aligned}
$$

where h, F and $\text{Cov}(e_t) = R$ and $\text{Cov}(u_t) = gg' = G$ are parameters. To adapt our model to this form, we use the following variable definitions

$$x_t = (z_{t-1}, z_t, f_t), \qquad F = (\beta, -\beta, \gamma),$$

$$h = \begin{pmatrix} 0 & 1 & 0 \\ 0 & \phi & 0 \\ 0 & 0 & \kappa \end{pmatrix}, \qquad g = \begin{pmatrix} 0 & 0 & 0 \\ 0 & \sigma_\epsilon & 0 \\ 0 & 0 & \sigma_f \end{pmatrix}.$$

The Kalman filtering recursions are as follows

$$\begin{aligned} \hat{\hat{x}}_t &= F\hat{x_{t-1}}, \\ \hat{\hat{P}}_t &= F'\hat{P_{t-1}}F + gg', \\ C_t &= h'\hat{\hat{P}}h + R, \\ \hat{e_t} &= y_t - \mu - h'\hat{\hat{x}}_t, \\ K &= \hat{\hat{P}}_t h(h'\hat{\hat{P}}_t h + R)^{-1}, \\ \hat{x_t} &= \hat{\hat{x}}_t + K\hat{e}_t, \\ \hat{P}_t &= \hat{\hat{P}}_t - Kh'\hat{\hat{P}}_t. \end{aligned}$$

The likelihood function can be computed $l = \sum_t -.5 \log |C_t| - .5\hat{e}_t{}' C_t^{-1} \hat{e}_t$. We employ a Metropolis Hastings sampler to collect random draws from the posterior distributions under a uniform prior on the parameter space.

2.3.1 Identification

The multivariate model specification in (5) is, as written, not identified. To see why this is the case, suppose we change σ_ϵ^2 by a factor of two as well as multiplying β by one half. The likelihood function will remain unchanged as a result of having re-scaled the parameters. Thus, we really cannot separately identify all the factor loadings at the same time as the factor variances. Also note that this re-scaling of the parameters leave the R^2 in equation (8) unchanged since β^2 and σ_ϵ^2 enter as a product in (8).

To facilitate estimation, therefore, we need to impose parametric constraints. There are two options. We can constrain β's and γ's in some way, for example by setting $\sum_i^N \beta_i = 1$ and $\sum_i^N \gamma_i$ across the N assets. This leaves σ_ϵ and σ_f as free parameters. A second possibility is to set σ_f and σ_ϵ at some value. We choose the restrictions $\sigma_f = 1$ and $\sigma_\epsilon = \sqrt{(1-\phi^2)/(2-2\phi)}$ which will give $\mathrm{Var}(z_t) = 1$ unconditionally. Note that these restrictions do not in any way alter the possible set of posterior distributions for R^2's since the re-scaling between γ/σ_f and β/σ_ϵ is arbitrary. The parametric constraints only apply to the multivariate model in (5) and not the univariate Campbell-Shiller model.

3 Empirical results

Table 1 reports posterior means and standard deviations for the univariate model in eqns. (1)-(3). As can be seen from the table, estimates of the mean reversion coefficients, ϕ, range from 0.84 to 0.96 across portfolios. Estimates of σ_ϵ range from 0.05 to 0.22 while σ_η estimates range from 0.06 to 0.2. It is interesting to note that all parameter estimates appear to increase or decrease almost uniformly in the size dimension. The amount of total variation, and also predictable variation, thus appear to vary strongly across portfolios of different capitalization.

There is a fair amount of posterior variability in all estimated parameters. For example, the posterior standard deviation of the variation in the random walk component, σ_η, is typically a little less than half that of the posterior mean. This translate into a significant posterior uncertainty about what fraction of the variation in stock returns can be attributed to predictable variation. Underscoring this point, figure depicts the posterior $2.5, 50$ and 97.5 percentiles of the $R^2(T)$. The line depicting the 2.5% lower posterior percentile of the R^2 distribution appears visually indistinguishable from the horizontal axis in the plots. This indicates that there well might be an insignificant amount of predictable variation in the data. On the other hand, the upper 97.5 percentile of the distributions in some cases (small portfolios) exceed 40% indicating that there indeed *could* be a significant amount of predictable variation in the returns. We conclude, therefore, that posterior distributions obtained for the respective parameters in the univariate Campbell-Shiller model are not informative enough to render a sharp conclusion about the amount of predictable variation in annual stock returns.

To gauge the statistical significance of the predictability in the univariate models, figure 2 plots the point-wise percentiles of the posterior distributions of $R^2(T)$'s. As can be seen from the plots, in all cases the lower 2.5-10% of the posterior mass of the $R^2(T)$'s near zero. In fact, in the posterior percentile plots the 2.5,5 and 10% percentile visually appear very close to zero. In other words, the analysis put significant posterior mass on the possibility that there is economically insignificant amounts of predictability in the data. Off course, there is also some posterior mass on the alternative - R^2 could likely be in the plus ten percent range at a five year forecasting horizon.

In assessing the reason why our analysis fail to detect significant amounts of predictability, figure 3 depicts empirical histograms of the posterior parameter draws for the middle sized/ BM portfolio. The others are have identically distributed parameters. The posteriors histograms indicate that ϕ has significant posterior mass near unity, and σ_ϵ has significant mass near zero. Both observations indicate predictable variation is statistically insignificant. On the other hand, σ_η also has mass near zero. Since all variation in the stock price is forecastable when σ_η is zero, the estimated model parameters also place significant posterior mass on the possibility of a large predictable component.

Table 1: Parameter Estimates for the Univariate Model

The table reports posterior means and standard deviations (in parenthesis), for parameters in the latent fundamentals model for Fama-French 5x5 size/book sorted portfolios. Results are based on annual portfolio returns from 1927 -2004.

	Low	2	3	4	High
			ϕ		
Small	0.84	0.86	0.88	0.89	0.89
	(0.11)	(0.09)	(0.09)	(0.09)	(0.09)
2	0.91	0.94	0.94	0.93	0.90
	(0.07)	(0.06)	(0.05)	(0.05)	(0.07)
3	0.93	0.95	0.95	0.94	0.91
	(0.05)	(0.04)	(0.04)	(0.04)	(0.06)
4	0.95	0.95	0.94	0.95	0.92
	(0.04)	(0.04)	(0.04)	(0.04)	(0.05)
Large	0.95	0.96	0.94	0.93	
	(0.04)	(0.03)	(0.04)	(0.05)	
			σ_η		
Small	0.22	0.21	0.17	0.16	0.16
	(0.10)	(0.09)	(0.09)	(0.08)	(0.08)
2	0.16	0.14	0.13	0.14	0.16
	(0.08)	(0.06)	(0.06)	(0.07)	(0.07)
3	0.14	0.12	0.11	0.12	0.16
	(0.07)	(0.05)	(0.05)	(0.05)	(0.07)
4	0.10	0.10	0.11	0.12	0.16
	(0.04)	(0.04)	(0.05)	(0.06)	(0.07)
Large	0.07	0.06	0.09	0.11	
	(0.03)	(0.03)	(0.04)	(0.05)	
			σ_ϵ		
Small	0.20	0.16	0.18	0.17	0.17
	(0.11)	(0.09)	(0.09)	(0.09)	(0.09)
2	0.16	0.14	0.13	0.15	0.14
	(0.08)	(0.07)	(0.06)	(0.07)	(0.07)
3	0.15	0.12	0.11	0.11	0.15
	(0.07)	(0.06)	(0.06)	(0.06)	(0.08)
4	0.09	0.09	0.10	0.12	0.15
	(0.05)	(0.05)	(0.05)	(0.06)	(0.08)
Large	0.06	0.05	0.08	0.10	
	(0.04)	(0.03)	(0.05)	(0.05)	

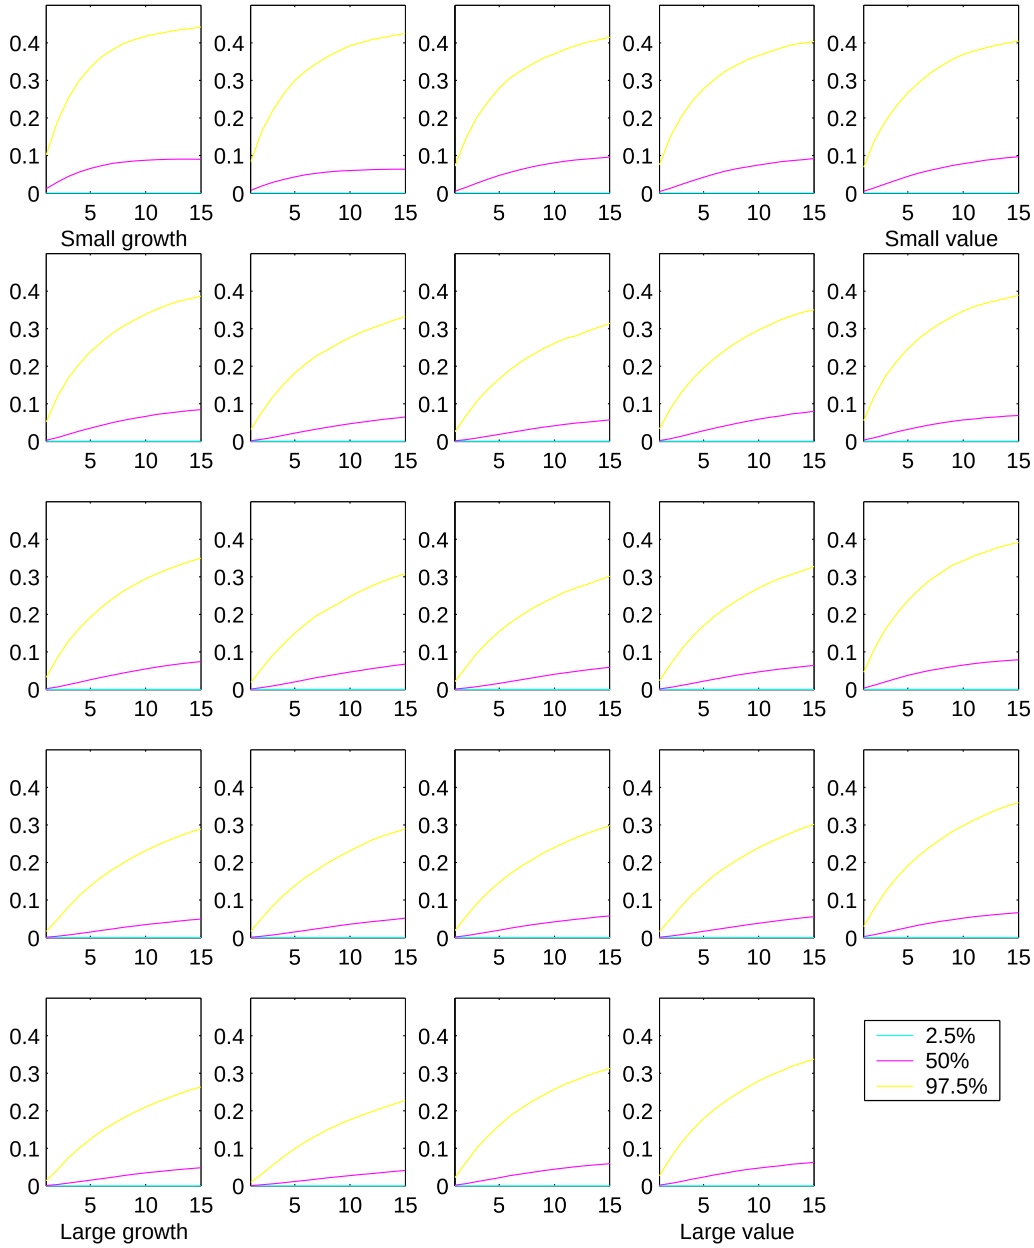

Figure 2: Pointwise posterior quantiles of $R^2(T)$.

Table 2: Univariate Model R^2

	Low	2	3	4	High
			1 year		
Small	0.01	0.01	0.01	0.00	0.00
2	0.00	0.00	0.00	0.00	0.00
3	0.00	0.00	0.00	0.00	0.00
4	0.00	0.00	0.00	0.00	0.00
Large	0.00	0.00	0.00	0.00	
			3 year		
Small	0.04	0.03	0.03	0.02	0.02
2	0.02	0.01	0.01	0.01	0.02
3	0.01	0.01	0.01	0.01	0.02
4	0.01	0.01	0.01	0.01	0.01
Large	0.01	0.01	0.01	0.01	
			5 year		
Small	0.07	0.04	0.05	0.04	0.04
2	0.04	0.02	0.02	0.03	0.03
3	0.03	0.02	0.02	0.02	0.04
4	0.02	0.02	0.02	0.02	0.03
Large	0.02	0.01	0.02	0.02	
			10 year		
Small	0.09	0.06	0.08	0.07	0.08
2	0.07	0.05	0.04	0.06	0.06
3	0.05	0.05	0.04	0.05	0.07
4	0.03	0.04	0.04	0.04	0.05
Large	0.03	0.03	0.04	0.05	
			30 year		
Small	0.08	0.06	0.10	0.10	0.10
2	0.10	0.09	0.08	0.11	0.07
3	0.10	0.10	0.09	0.09	0.09
4	0.07	0.08	0.08	0.08	0.08
Large	0.07	0.06	0.08	0.08	

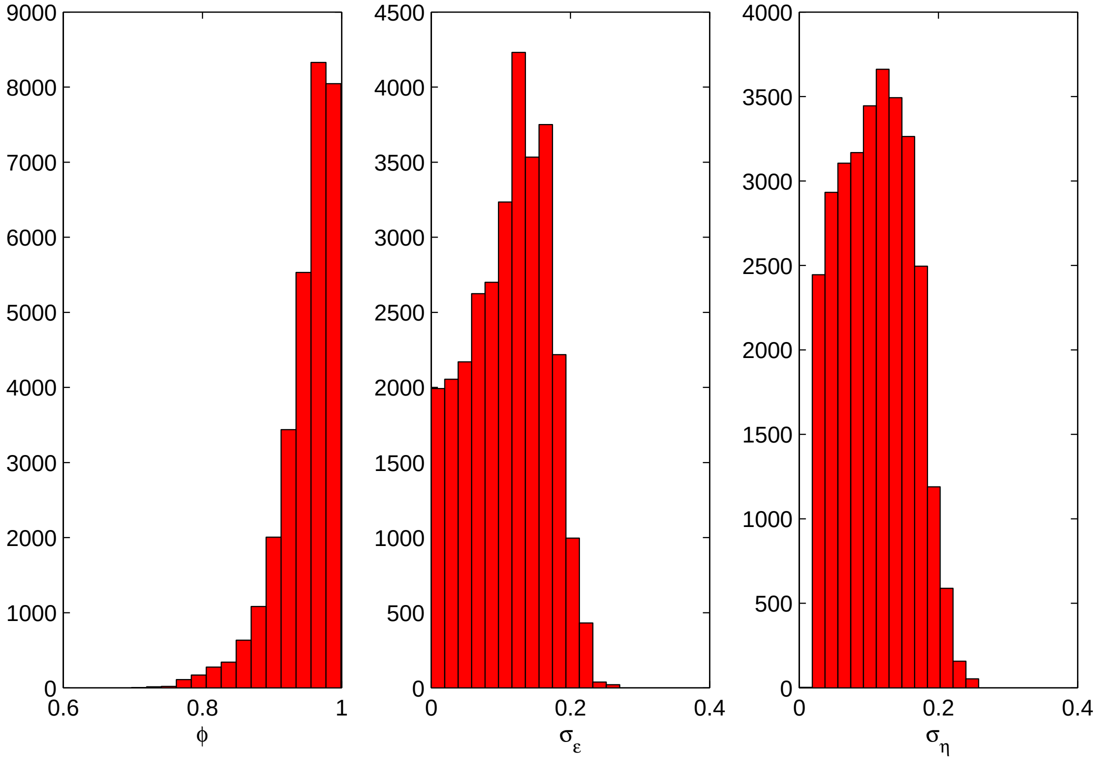

Figure 3: Posterior histograms for univariate model parameters.

3.1 Results from Multivariate Models

Table 3 reports the posterior estimates of the parameters in the multivariate model, eqn. (5)-(7). The results are based on the Fama-French five-by-five size and book/market sorted portfolios[2]. The results show several interesting features. First, note that the mean of the posterior for ϕ is much smaller than for the univariate models, indicating that the mean reversion is quicker in the multivariate data-set. The short term factor f_t has an autocorrelation of $\kappa = -.074$ with posterior standard deviation 0.167. Thus, the posterior distribution of κ is close to zero, indicating that this factor is picking up almost no predictable variation in the returns. In conclusion, almost all of the predictable variation is found to come from the long-term factor z_t.

[2]We exclude the 25th portfolio containing the largest firms with the highest book/ market values due to missing observations.

Table 3: Multivariate Model Parameter Estimates

	ϕ	κ	ρ
	0.838 (0.088)	-0.074 (0.167)	0.899 (0.030)

	Low	2	3	4	High
			β		
Small	0.273	0.315	0.311	0.319	0.317
2	0.245	0.230	0.240	0.255	0.276
3	0.204	0.201	0.201	0.210	0.250
4	0.124	0.167	0.172	0.190	0.204
Large	0.076	0.084	0.093	0.127	
			γ		
Small	-0.040	0.000	0.018	0.050	0.054
2	-0.068	0.005	0.035	0.074	0.087
3	-0.069	0.021	0.053	0.079	0.146
4	-0.040	0.020	0.080	0.106	0.146
Large	-0.005	0.049	0.097	0.145	
			μ		
Small	-0.013	0.056	0.099	0.130	0.141
2	0.048	0.101	0.119	0.122	0.125
3	0.069	0.104	0.112	0.117	0.112
4	0.081	0.090	0.107	0.110	0.105
Large	0.078	0.080	0.086	0.087	

Table 4: Multivariate Model R^2

	Low	2	3	4	High
			1 year		
Small	0.06	0.11	0.12	0.13	0.13
2	0.11	0.11	0.11	0.10	0.10
3	0.09	0.09	0.09	0.09	0.07
4	0.05	0.07	0.06	0.06	0.05
Large	0.02	0.02	0.02	0.03	
			3 year		
Small	0.10	0.17	0.20	0.20	0.20
2	0.17	0.17	0.18	0.16	0.16
3	0.15	0.14	0.13	0.13	0.10
4	0.08	0.11	0.09	0.09	0.06
Large	0.03	0.03	0.02	0.03	
			5 year		
Small	0.10	0.18	0.20	0.21	0.20
2	0.18	0.17	0.18	0.16	0.16
3	0.15	0.14	0.14	0.13	0.10
4	0.08	0.12	0.09	0.09	0.06
Large	0.03	0.03	0.02	0.03	
			7 year		
Small	0.09	0.17	0.19	0.20	0.20
2	0.17	0.16	0.17	0.15	0.15
3	0.14	0.13	0.13	0.12	0.09
4	0.07	0.11	0.08	0.08	0.06
Large	0.03	0.03	0.01	0.02	
			10 year		
Small	0.08	0.15	0.17	0.18	0.18
2	0.15	0.14	0.15	0.13	0.13
3	0.12	0.11	0.11	0.11	0.08
4	0.06	0.09	0.07	0.07	0.05
Large	0.02	0.02	0.01	0.02	

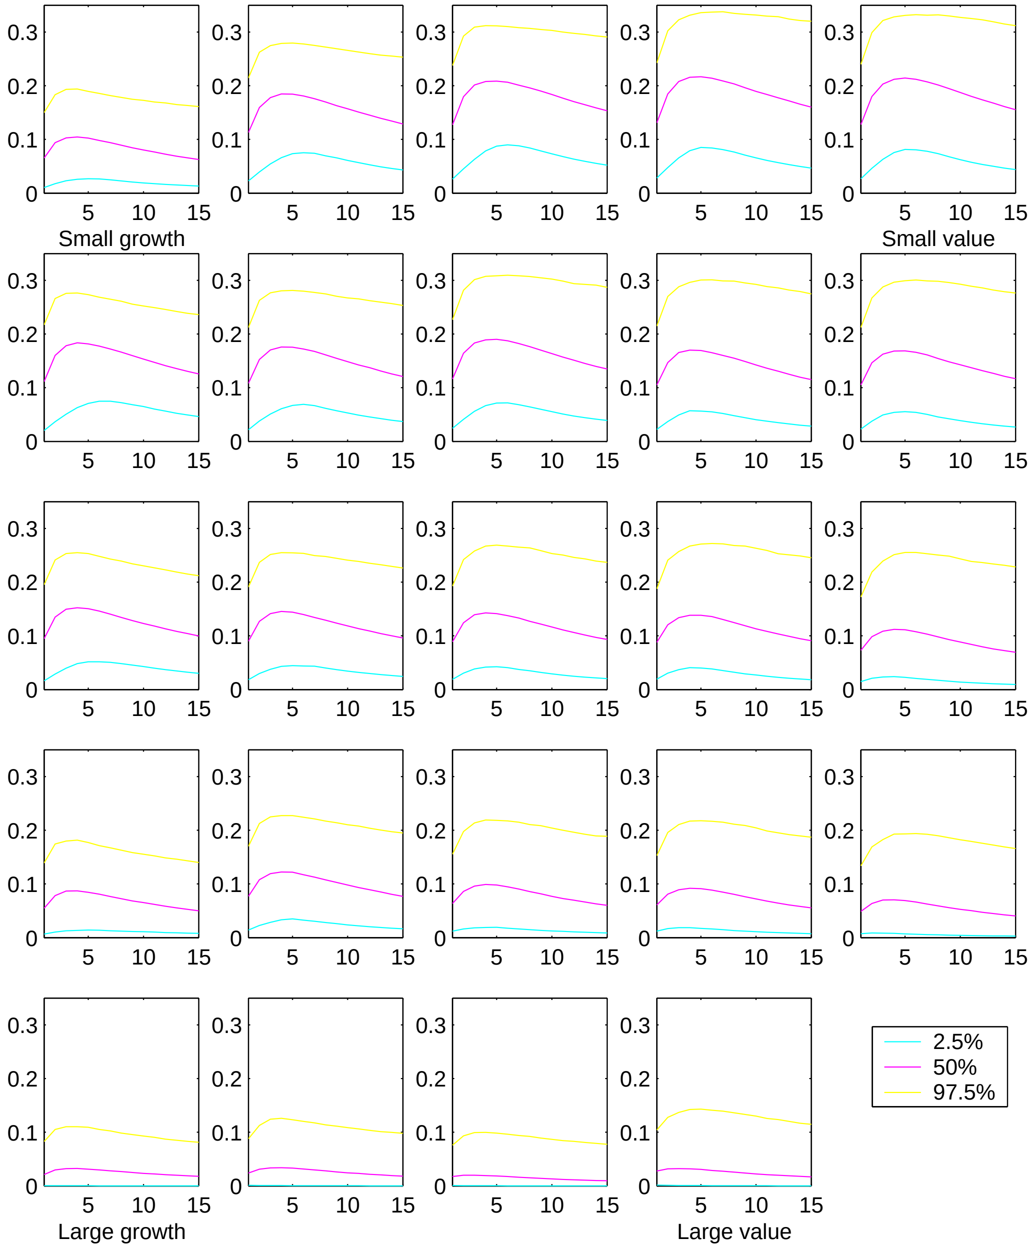

Figure 4: Pointwise posterior quantiles of $R^2(T)$. Multivariate model.

Table 3 shows some interesting differences in the predictable return components along the size and book/market dimensions. The factor sensitivities, β, associated with the long-term factor z_t, are strongly decreasing with size. Small firm portfolios have β's roughly three times that of large firm portfolios. This would indicate that the predictable variation in small firm portfolios is much more pronounced than in large firm portfolios. The table also shows that posterior means for β increase along the value/growth dimension, as portfolios of high book/market (value) firms generally have higher β's. Overall, therefore, we expect predictability to be stronger in small, value portfolios.

Table 4 illustrates that this is the case by tabulating the posterior median R^2 as computed using (8). This table reveals that R^2's are generally much larger for the small firm portfolios than for large, and generally somewhat larger for high B/M portfolios. Table 4 illustrates that small, value portfolios have R^2's that peak at around 20% at the three to seven year forecasting horizon. We find that large firms have negligible R^2's at almost all forecasting horizons, suggesting that predictability, if any, is largely a small firm phenomena.

Figure 4 depicts the 2.5%, 50% and 97.5% posterior percentiles of the R^2's. The figure confirms the conclusion that predictability is minimal in large firm portfolios, as the lower 2.5% percentile is visually indistinguishable from the horizontal axis. In most cases also the upper 97.5 percentiles peak at values less than 15%, suggesting that long term mean reversion is an unlikely occurrence in large firm stock. For small value firms, the situation is different. These portfolios have point-wise 2.5 percentiles that peak at values in the 7-8% range. We interpret this as statistically significant[3]. Fama and French (1988) find that small firm portfolios have larger long run autocorrelations, however, the difference between large and small firms is less substantial than results reported in table 4 and figure 4.

[3]The posterior probability that $R^2 < \epsilon$ for any ϵ less than a particular portfolios 2.5 percentile R^2 is less than 2.5%.

The "short-term factor," f_t, does not contribute a significant amount of predictable variation in the observed returns. This is evidenced by the fact that κ has a significant posterior mass both above and below zero. When $\kappa = 0$, f_t is a random walk which again implies that f_t does not contribute to explaining any predictable variation in the observed returns.

Are the predictable patterns in stock return data found here consistent with equilibrium? Recent advances in the literature on consumption based asset pricing models with time-varying risks and expected returns suggests that the cross sectional premiums found in book and size sorted portfolios may well be consistent with rational pricing. For example, Bansal, Dittmar and Lundblad (2003,2004) find that differences in so-called cash flow betas explain cross-sectional differences in average rates of return across size and book-to-market assets. The cash flow beta is defined as the covariance of a firms' cash flow NPV to consumption innovations.

The results reported here are consistent with a wide range of risk based equilibrium explanations. Indeed, a number of models that prescribe a cointegration relationship between stock prices and the aggregate dividends, consumption, or some other variable with a long-term random walk component will generate predictability patterns which are broadly consistent with those found here. For example Bansal, Dittmar and Kiku (2005) find that small firm dividends have larger co-integration parameters with aggregate consumption. Their model generates large R^2's for small firms at short horizons, and R^2's ranging from 36% to 59% at long (ten year) horizons. The fact that their numbers significantly exceed the ones presented here is natural because their model uses a larger information set to predict the returns.

We have concluded that κ is "insignificant," and the factor f_t mimics a random walk. It is interesting in this context that the factor loading relating to f_t, γ, are almost monotonically increasing in the book/market deciles. For the lowest B/M portfolios, the γ's are negative near $-.05$, while the high B/M portfolios have γ's ranging from 0.054 to 0.146. It is reasonable to interpret the differences in the γ's as evidence of a permanent difference in covariance structure, and hence risk, between portfolios of different characteristics.

3.2 Robustness

It is well known that the infamous U shaped autocorrelation function found empirically by Fama and French (1988) is not robust across all sampling periods. In particular, Fama and French themselves point out that the pattern is significantly less pronounced in post war data. It is widely believed that the highly volatility depression era is not representative for the subsequent sample. It is also easy to criticize the model presented here as it does not account for stochastic, time-varying volatility. Volatility is particulary high prior to the second world war. Table 5 reports parameter estimates computed over the 1946 to 2003 period. The table reveals that the estimated parameters are quite similar to those in table 3. Figure 3.2 shows that small firm R^2 are significant and median posterior R^2 peak at around 20% for small value portfolios at the ten year horizon. The lower 2.5 percentiles peak at around 10% for these portfolios. Overall the R^2's increase relative to the ones obtained using the whole sample. The estimates seem to imply somewhat larger importance of the "short term factor" f_t, reflected by a larger posterior mean of κ in absolute value. Overall, the evidence in table 5 is largely consistent with those obtained in the full sample in table 3.

In order to further assess the effects of time-varying volatility on our conclusions, we ran an additional estimation under the assumption that the local volatility of returns are generated according to a GARCH(1,1) model. The standardized returns were then used to re-estimate the model parameters. The estimates were quite similar to those in table 5 and thus not reported. Figure 3.2 shows posterior R^2's resulting from this exercise. The evidence is still suggestive of significant amounts of predictability in small value stocks, although somewhat weaker than for the full sample.

Table 5: Parameter Estimates using Post-War data

	ϕ	κ	ρ
	0.917 (0.057)	-0.183 (0.184)	0.915 (0.015)

	Low	2	3	4	High
			β		
Small	0.259	0.269	0.242	0.234	0.251
2	0.204	0.186	0.189	0.191	0.194
3	0.142	0.154	0.149	0.163	0.192
4	0.095	0.109	0.120	0.131	0.145
Large	0.023	0.044	0.044	0.069	
			γ		
Small	-0.150	-0.044	-0.010	0.030	0.035
2	-0.119	-0.029	0.013	0.057	0.061
3	-0.124	-0.006	0.039	0.067	0.083
4	-0.110	0.006	0.041	0.066	0.062
Large	-0.067	-0.006	0.004	0.068	
			μ		
Small	0.052	0.116	0.129	0.154	0.169
2	0.073	0.118	0.142	0.147	0.164
3	0.091	0.127	0.129	0.146	0.154
4	0.104	0.112	0.140	0.138	0.147
Large	0.105	0.109	0.125	0.124	

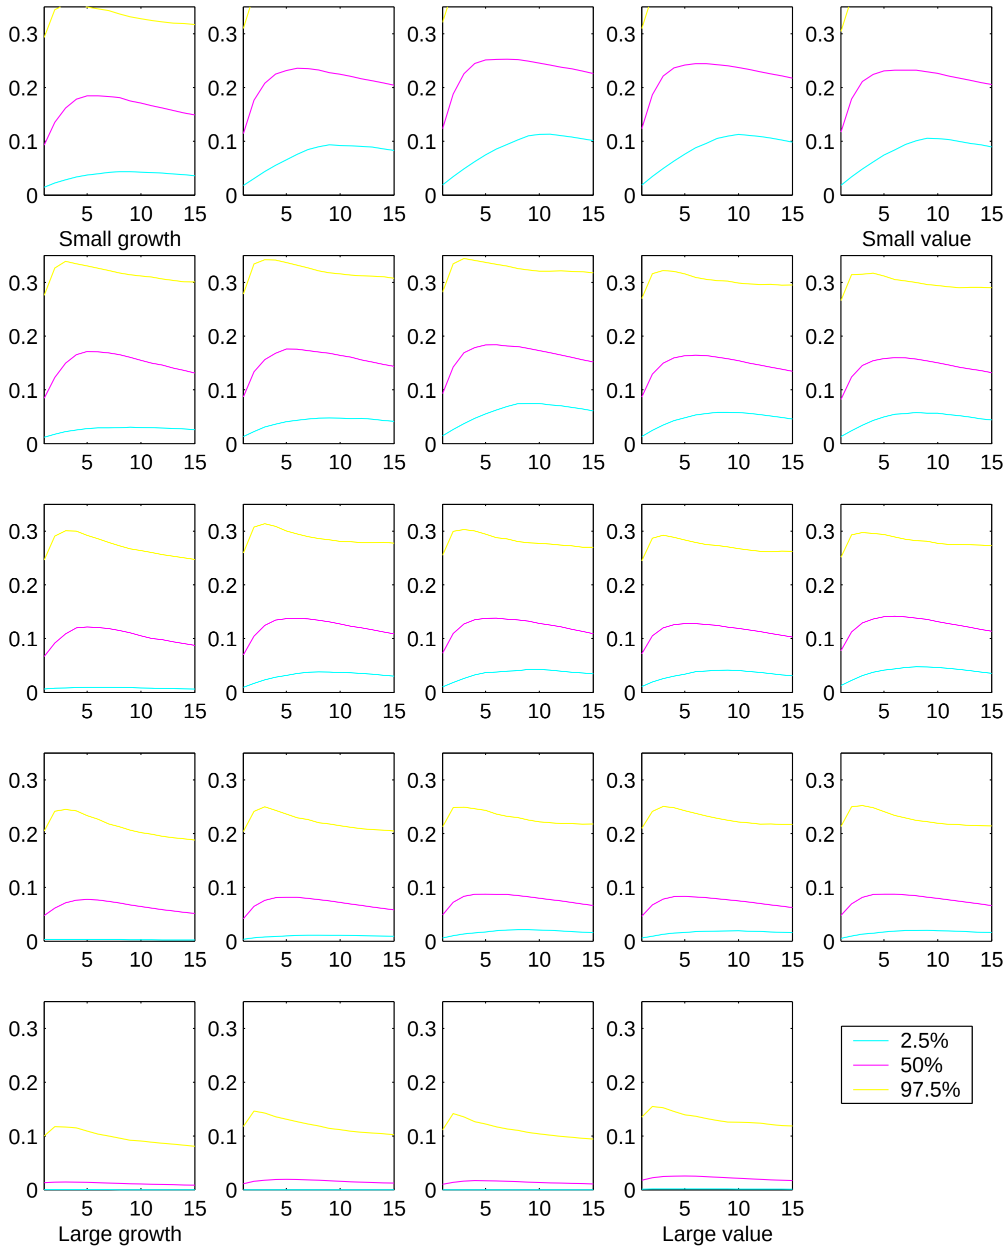

Figure 5: Posterior $R^2(T)$ quantiles for the multivariate model using post-war data.

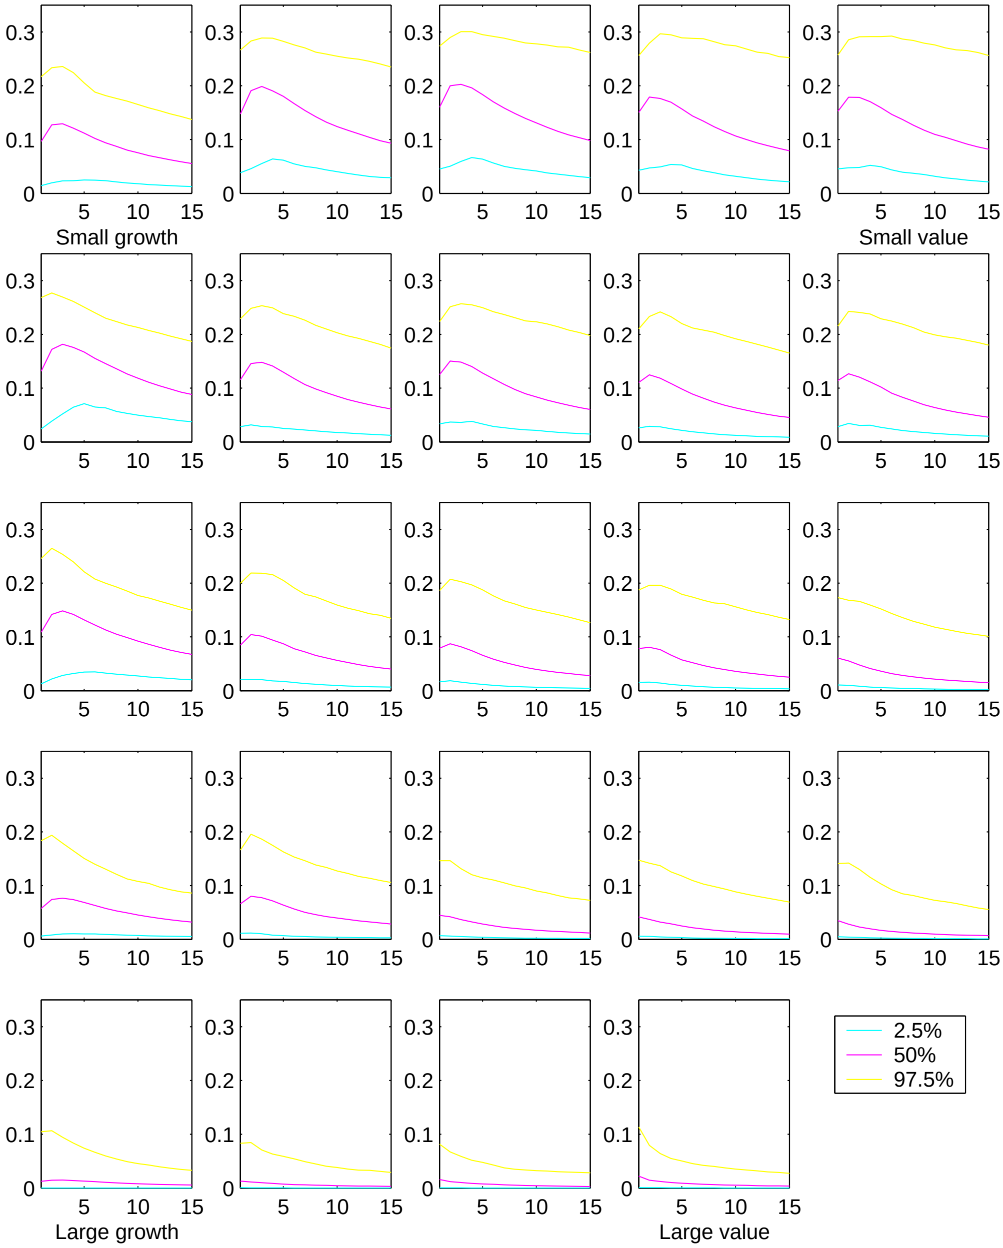

Figure 6: Posterior $R^2(T)$ quantiles for the multivariate model using data standardized by GARCH(1,1) time-varying volatility.

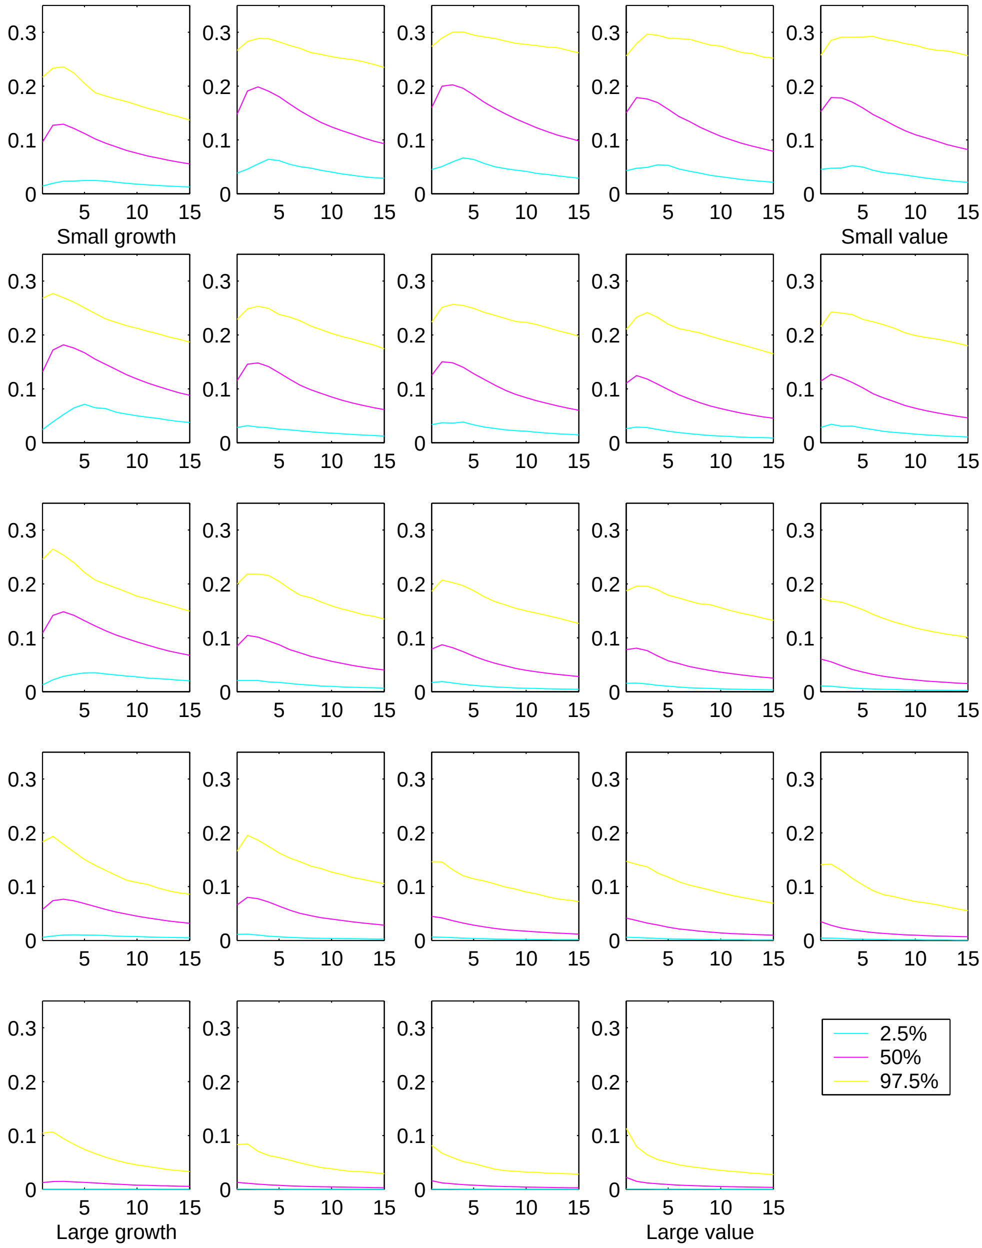

Figure 7: Posterior $R^2(T)$ quantiles for the multivariate model using data standardized by GARCH(1,1) time-varying volatility.

3.3 Latent Factors

A particular advantage of the Kalman filtering approach to estimate these models is that we recover filtered estimates of the unobserved state variables in the model. Figure 8 plots the filtered values at the posterior mean of the parameters as well as the model implied expected rates of return for one, three and five year forecasting horizons. Most of the variation in expected rates of return is due to variation in z_t. This is visually evident in the figure. The model generates significant time-variation in expected rates of return. This is particularly true at long forecasting horizons. For example, at the end of 2003, the forecast is for small firms to beat large by 20% cumulative return for the 2004 to 2009 period, corresponding to about 4% annually.

Another interesting aspect of the models considered in this paper is that they provide a different mechanism to compute expected returns than autocorrelations. Forecasts based on autocorrelations are formed with a small number of observations preceding the date of the forecast. Within the models considered here expected returns are formed using filtered values of the unobserved state-variables which utilizes the whole sample prior to the date of the forecast. Lewellen (2001) shows that predictability is significant when running regressions of twelve to eighteen month returns onto one to five year lagged returns. The fact that the longer lag structure improves the forecasting performance is consistent with the models and results presented here.

4 Conclusion

This paper has studied dynamic latent component models of stock price data. We find that a significant fraction of variation in small firm portfolios is due to predictability. Conversely, large firm stock portfolios contain insignificant predictable components.

There are multiple ways in which the models considered here could be generalized. It is straightforward to include covariate information to increase the amount of predictability. While we have demonstrated that pre-filtering the data with time

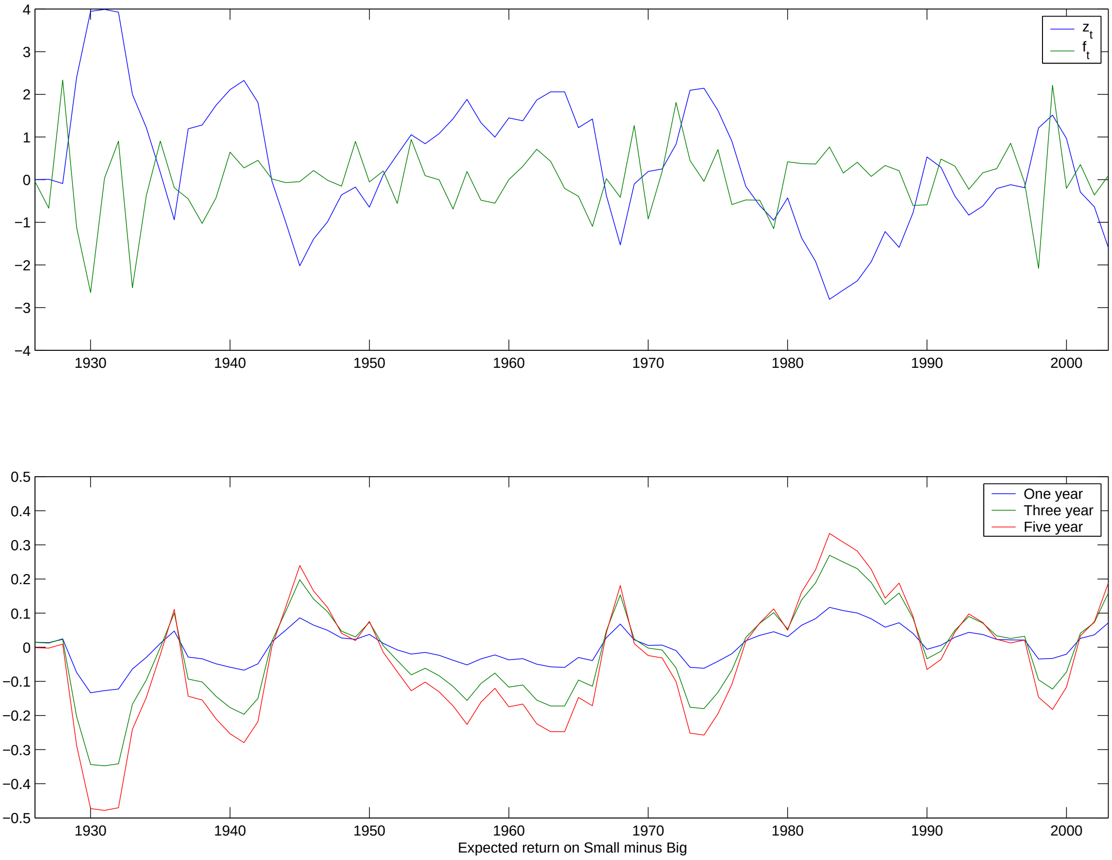

Figure 8: Filtered factors (above) and expected rates of return on small minus big firm portfolios (below).

varying GARCH volatiltiy does not alter the conclusions, it is possible to include time-varying volatility as part of the model. This is a particularly relevant extension in the context of higher frequency returns data. Higher frequency data will not allow us to identify the long term factor, z_t, with much greater precision, however. Higher frequency data will on the other hand allow the identification of a short term factor such as f_t, assuming that such a factor is relevant. The addition of stochastic volatility is complicated in the high dimensional setting such as here because the likelihood function must be computed by numerical integration which is typically done though Monte-Carlo.

Our analysis complements the many recent studies of dynamic equilibrium models which generate time-variation in expected rates of return. In particular, our results are consistent with risk premia driven differences in conditional and unconditional expected rates of return across size and book/market assets. That said, this analysis is not a test of equilibrium and as such and the results may also be consistent with irrational pricing. A natural extension of this study is to seek to identify risk-based sources of variation in the long term expected rates of return.

References

Bansal, Ravi, Robert F. Dittmar, and Dana Kiku, 2005, Long Run Risks and Equity Returns, *Working paper, Duke.*

Bansal, Ravi, Robert F. Dittmar, and Christian Lundblad, 2004, Consumption, Dividends, and the Cross-section of Equity Returns, *Journal of Finance* forthcoming.

Bansal, Ravi, and Amir Yaron, 1997, Risks for the Long Run: A Potential Resolution of Asset Pricin Puzzles, *Journal of Finance* 59, 1481–1509.

Bekaert, Geert, and Campbell R. Harvey, 1995, Foreign Speculators and Emerging Equity Markets, *Journal of Finance* 55, 565–613.

Brandt, Michael W., and Qiang Kang, 2004, On the Relationship between Conditional Mean and Volatility of Stock Returns: A latent VAR approach, *Journal of Finance* 72, 217–257.

Campbell, John Y., and Robert J. Shiller, 1987, Cointegration and Tests of Present Value Models, *Journal of Political Economy* 95, 1062–1088.

Campbell, John Y., and Robert J. Shiller, 1988, Stock Prices, Earnings and Expected Dividends, *Journal of Finance* 43, 661–676.

Fama, Eugene F., 1991, Efficient Capital Markets II, *Journal of Finance* 46, 1575–1617.

Fama, Eugene F., and Kenneth R. French, 1988, Permanent and Temporary Components of Stock Returns, *Journal of Political Economy* 96, 246–273.

Hansen, Lars P., John Heaton, and Nan Li, 2004, Consumption Strikes Back?, .

Johannes, Michael, Nick Polson, and John Stroud, 2004, Sequential Optimal Portfolio Performance: Market and Volatility Timing, *Working Paper, Columbia.*

Lamoureux, Christopher G., and Goufu Zhou, 1996, Temporary Components of Stock Returns: What Do the Data Tell Us?, *Review of Financial Studies* 9, 1033–1059.

Lewellen, Jonathan, 2001, Temporary movements in stock prices, *Working paper, MIT.*

Poterba, James, and Lawrence Summers, 1988, Mean Reversion in Stock Prices: Evidence and Implications, *Journal of Financial Economics* 22, 27–59.

Richardson, Matthiew, 1993, Temporary Components of Stock Prices: A Sceptic's View, *Journal of Business and Economic Statistics* 11, 199–207.