

A Dynamic Test of Conditional Factor Models*

Daniele Bianchi[†]

Abstract

I use Bayesian tools to dynamically test conditional factor pricing models from the point of view of an investor who recognizes that parameters are uncertain, time-varying, and predictors are an imperfect proxy for macroeconomic and firms-specific news. Time-varying alphas, betas and idiosyncratic risks are jointly estimated in a single-step together with no-arbitrage restrictions. The test can be applied for a single asset or jointly across portfolios. As empirical application, I estimate over fifty years of post-war monthly data a conditional version of the CAPM and multi-factor models on size, book-to-market and momentum deciles portfolios. I show that once the dynamic and uncertain nature of the portfolios returns generating process is fully acknowledged, the null that conditional factor models hold is not sensibly rejected both in the time series and in the cross-section.

Keywords: Conditional Factor Models, Pricing Anomalies, Stochastic Betas, Bayesian Econometrics

JEL codes: G12, E44, C11

*This version: September 30, 2014, **Comments are welcome.** I am grateful to Carlos M. Carvalho, Roberto Casarin, Massimo Guidolin, John M. Maheu, Lasse H. Pedersen and Rodney W. Strachan for their useful comments and suggestions. I also thank seminar participants at the University of Warwick, Bocconi University, University of Venice Cá Foscari, and the 8th Bayesian econometrics workshop RCEA. The paper is based on my dissertation at the Department of Finance, Bocconi University. All remaining errors are mine.

[†]Warwick Business School, University of Warwick, Coventry, CV4 7AL, UK. Daniele.Bianchi@wbs.ac.uk

1 Introduction

Traditional empirical methodologies to test mean-variance efficiency in factor asset pricing models assume either factor loadings are constant, parameters are observable, or eventual time variation in risk exposures can be exactly characterized by a set of known predictors. However, a long history of evidence shows that factor loadings are not constant over time, parameters are uncertain, and risk exposures respond to macroeconomic and firm-specific news which are, at least for the most part, unpredictable.¹

I introduce a methodology to jointly estimate time-varying alphas, betas, idiosyncratic risks and factors risk premia. Such a joint estimation scheme sidesteps the crucial issues of traditional two-steps regression-based methods (see e.g., Petersen 2009 and Kan, Robotti, and Shanken 2013). I develop estimators condition on available information about portfolios returns, but realizing knowledge is limited in two key aspects. First, the *true* parameters of the returns generating process are not observable. While acknowledging that betas and idiosyncratic risks vary considerably over time, such uncertainty on structural parameters indeed persists even after observing 50 years of portfolios returns. Second, observable predictors used to forecast the betas and the factors risk premia deliver only an imperfect proxy for macroeconomic and company-specific news, implying a stochastic dynamics. I use Bayesian tools, which, by construction, helps generate posterior distributions of virtually any function of the structural parameters/statistics of factor models.

Based on such estimation procedure, I propose a methodology to dynamically test conditional factor models both in the time series and in the cross-section. Under the null of a correctly specified factor model, an asset's unexplained return should not be statistically significant over time after controlling for the asset's exposures to sources of systematic risks. Such null hypothesis can be seen as a linear restriction on the dynamics of a more general factor pricing model. I develop a method to compute the dynamic posterior probabilities of linear restrictions on assets' excess return using the output of the joint estimation scheme and the principle of the Savage-Dickey density (SDD) ratio (see Verdinelli and Wasserman 1995 and

¹See among others, Fama and French (1997), Harvey (2001), Jostova and Philipov (2005), Lewellen and Nagel (2006), Ang and Chen (2007), Nardari and Scruggs (2007), Adrian and Franzoni (2009), Petersen (2009), Ang and Kristensen (2012) and Bianchi, Guidolin, and Ravazzolo (2014).

Koop, Leon-Gonzalez, and Strachan 2010). Major advantages of this methodology are that; (1) fully integrates parameter uncertainty, (2) can be applied to both single assets and across portfolios, and (3) is recursive in nature as built conditional upon the information available at a given time t . More prominently, the dynamic nature of the test allows to compute the posterior probability that the null of a factor model holds at time t , but does not necessarily holds at other times.

Traditional empirical methodology to test factor pricing models implement rolling window estimates of asset betas from time series regressions, and use them to estimate risk premia via cross-sectional regressions (e.g. Fama and MacBeth 1973). Unless additional assumptions are introduced, time-varying factor loadings can distort standard tests which makes most of the inferential statements commonly made invalid (see Pagan 1984, Ang and Chen 2007, Petersen 2009 and Kan et al. 2013). Even though asymptotic adjustments can be introduced, it remains unclear if such bias-corrected inference is indeed reliable in finite samples, especially given the sample size commonly used in testing conditional factor pricing models. I take important steps beyond this approach by developing an exact finite-sample testing framework for conditional factor pricing models.

This paper builds on a literature advocating the use of Bayesian methods to estimate time-varying risk exposures and risk premia, and more generally test asset pricing models, such as McCulloch and Rossi (1991), Geweke and Zhou (1996), Jostova and Philipov (2005), Ang and Chen (2007) and Nardari and Scruggs (2007). In particular, Geweke and Zhou (1996) show how to obtain exact posterior distributions for functions of interest in factor pricing models, while Ang and Chen (2007) show that under the null of time-varying betas, standard OLS inference produces inconsistent estimates. In the spirit of McCulloch and Rossi (1991) and Geweke and Zhou (1996), I provide an exact finite-sample testing framework for APT-like factor models.

My work extends this literature in several important ways. I provide a joint estimation framework for time-varying betas, idiosyncratic risks and risk premia without having to instrument risk exposures and premia with observable predictors. For example, Geweke and Zhou (1996) provide a single-step Bayesian estimation setting to sidestep the drawbacks of traditional regression-based tests on factor pricing models. This turns out to be a special case in this paper where betas and idiosyncratic risks are considered constant over time. The economet-

ric framework parallels that in Bianchi et al. (2014), who focus on macro-based factor pricing models and investigate the effect of having discrete breaks in the dynamics of structural parameters. In addition, I provide a dynamic method to recursively test the null of a conditional factor model both in the time series and in the cross-section in a unified setting, which the earlier literature did not provide.

Empirically, the paper is focused on conditional versions of some of the most common factor pricing models. These are tested both in the time series and in the cross-section, both on single portfolios and jointly across assets. In addition to the basic conditional CAPM, I study the well-known three-factor model of Fama and French (1993), together with the four-factor model proposed in Charhart (1997). In general, their size, value and momentum factors are assumed to capture investors' expectations on business cycle effects (see e.g. Liew and Vassalou 2000, Cochrane 2001, Vassalou 2003 and Campbell and Diebold 2009). The main empirical analysis uses the standard 25 size and book-to-market portfolios of Fama and French (1993) and ten momentum deciles portfolios as test assets. Momentum portfolios are included to mitigate the strong common structure which arguably characterize the Fama-French portfolios (see Lewellen, Nagel, and Shanken 2010 for a detailed discussion). I show that, once the dynamic and uncertain nature of the portfolios returns generating process is fully acknowledged, time series test do not reject the null of factor pricing models. Interestingly, the results on the recursive cross-sectional test highlights the outperformance of the three-factor Fama-French model. Indeed, conditional specifications of both the CAPM and the four-factor model reject the null of no pricing error in the cross-section across the period 2001/2002 (i.e. 9/11, Financial scandals, Iraq war) and across the recent great financial crisis.

Consistent with previous literature, I show that conditional betas do vary considerably over time and differ across different portfolios. Time series averages exhibit the usual cross-sectional pattern. Small-cap stocks are riskier than large-cap stocks, stocks with high book-to-market ratios are riskier than those with low ratios, and past losers are on average riskier than past winners. The exposures of momentum-sorted portfolios to both *size* and *value* factors are not systematically different from zero across the sample. Idiosyncratic risks also vary considerably over time, and peaks around the period 2001/2002 (dot.com bubble burst, financial scandals, Gulf War II, 9/11 attacks), and across the recent financial crisis.

The remainder of the paper proceeds as follows. Section 2 lays out the empirical framework. Section 3 reports the data, prior calibration and reports the main empirical results. Section 4 discusses sources of instability in the betas and the role of idiosyncratic risks. Next, Section 5 presents further assessment of models performances, then Section 6 concludes. I leave to the appendix derivation details.

2 A Bayesian Framework for Conditional Factor Models

Conditional factor pricing models posit a linear relationship between excess returns on N assets and a set of $K \ll N$ common tradeable factors. In general, these factors are assumed to capture investors' expectations on business cycle effects (see e.g. Liew and Vassalou 2000, Cochrane 2001, Vassalou 2003 and Campbell and Diebold 2009). If we call the process for the risk factors $F_t = [F_{1,t}, \dots, F_{K,t}]'$ and $y_{i,t}$ the period excess return on asset or portfolio $i = 1, \dots, N$, computed as $y_{i,t} = [(P_{i,t} - P_{i,t-1} + D_{i,t})/P_{i,t-1}] - r_{f,t}$ where $P_{i,t}$ denotes the price of any asset or portfolio, $D_{i,t}$ any dividend or cash flow paid out by the asset, and $r_{f,t}$ the one-period interest rate, a typical conditional factor model can be written as:

$$y_{i,t} = \beta_{i0,t} + F_t' \beta_{i,t} + \sigma_{i,t} \epsilon_{i,t} \quad \epsilon_{i,t} \sim N(0, 1) \quad i = 1, \dots, N \quad (1)$$

where $\beta_{i,t} = [\beta_{i1,t}, \dots, \beta_{iK,t}]'$ the K -dimensional vector of asset specific exposures to risk factors. The conditional alphas $\beta_{i0,t}$ is often interpreted as pricing error in the time series of the i th portfolio return. Indeed, if $F_t = 0$, then (1) implies that $y_{i,t} = \beta_{i0,t} + \sigma_{i,t} \epsilon_{i,t}$. This violates mean-variance efficiency stating that in absence of any risk factor, the risk premium for the i th portfolio turns out to be different from zero. In the conditional version of Ross' (1976) APT, or in Merton (1973) inter-temporal CAPM (ICAPM), the pricing kernel, M_{t+1} , must be linearly dependent on the K -dimensional vector of risk factors. As such, under no-arbitrage opportunities, the fundamental pricing equation $E_t [M_{t+1} y_{i,t+1}] = 0$ implies that

$$E_t [y_{i,t+1}] \approx Var_t [F_{t+1}] \times \left(\frac{Cov_t [y_{i,t+1}, F_{t+1}]}{Var_t [F_{t+1}]} \right) = \gamma_t' \beta_{i,t} \quad (2)$$

namely, the expected excess returns of the i th asset/portfolio over the interval $[t, t+1]$, $E_t[y_{i,t+1}]$, is related to its “current” betas, $\beta_{i,t}$, and the factors risk premia, $\gamma'_t = [\gamma_{1,t}, \gamma_{2,t}, \dots, \gamma_{K,t}]$ (see Cochrane 2001 for more details). This no-arbitrage restriction is known to hold under a variety of alternative assumptions, and conditional on the information publicly available at time t (see Bossaerts and Green 1989 and Ferson and Harvey 1991). The equilibrium condition (2) does not imply any statistical significant intercept $\gamma_{0,t} \neq 0$. Indeed, in the absence of arbitrage all zero-beta assets should command a rate of return that equals the short-term rate, which is zero if no risk-less investable assets are considered.

2.1 Bayesian Estimation with Pricing Restrictions

The asset pricing theory leaves unspecified the exact nature of conditioning information underneath the dynamics of alphas, betas and idiosyncratic risks. Early approaches describe alphas and betas as linearly dependent in a set of observable state variables aimed at capturing broad economic conditions. These require the investors to exactly know the right state variables (e.g. Shanken 1990, Jagannathan and Wang 1996, Lettau and Ludvigson 2001 and Adrian and Franzoni 2009). Other approaches, such as French, Schwert, and Stambaugh (1987), Campbell and Voulteenahe (2004), Fama and French (2005), Lewellen and Nagel (2006) and Ang and Kristensen (2012), rely on a series of constant parameter models in the spirit of Fama and MacBeth (1973).

By assuming that a small set of instrumental variables approximate a potentially large amount of both macroeconomic and company-specific news might lead to misleading inference, as the factors betas become very sensitive to the choice of instruments (see e.g. Harvey 2001). Also, Ang and Chen (2007) show that constant parameter models induce to biased inference under the null of time varying alphas and betas. Even using high-frequency returns to approximate the dynamics of factors risk exposures, it is not clear how we can link those estimates to test for the cross-sectional implications implied by the beta representation in (2). Indeed, two-steps procedures such as the Fama and MacBeth (1973) are quite inefficient (see e.g. Petersen 2009 and Bianchi et al. 2014). In this paper, I characterize the relationship between excess returns, factors and risk premia, as well as the time-varying dynamics in factor loadings and idiosyncratic volatility as a state-space model where the linear factor model (1) and the

non-linear no-arbitrage restriction (2) are jointly considered;

$$y_{i,t} = \beta_{i0,t} + F_t' \beta_{i,t} + \sigma_{i,t} \epsilon_{i,t} \quad \epsilon_{i,t} \sim N(0, 1) \quad (3)$$

$$y_{i,t} = \gamma_{0,t} + \gamma_{1:K,t}' \beta_{i,t-1} + e_{i,t} \quad e_{i,t} \sim N(0, \omega^2) \quad (4)$$

where $E[\epsilon_{i,t}] = E[\epsilon_{i,t} F_{j,t}] = E[e_{i,t} \beta_{ij,t-1}] = 0$ for all $i = 1, \dots, N$ and $j = 1, \dots, K$. The error term $e_{i,t}$ is due to the fact that (4) represents a statistical approximation of the equilibrium condition (2). Time varying parameters $\beta_{ij,t}$ and σ_{it} are described as latent states without specifying ad-hoc conditioning information. The implicit assumption is that macroeconomic and company-specific news that affect risk exposures either cannot be exactly anticipated or arrive randomly (e.g. Jostova and Philipov 2005). Therefore, alphas, betas and idiosyncratic risks evolve for each $i = 1, \dots, N$ and $j = 0, \dots, K$ as

$$\beta_{ij,t} = (1 - \delta_{ij}) \overline{\beta_{ij}} + \delta_{ij} \beta_{ij,t-1} + \tau_{ij} \eta_{ij,t} \quad \eta_{ij,t} \sim N(0, 1) \quad (5)$$

$$\ln(\sigma_{i,t}^2) = (1 - \delta_{i\sigma}) \ln \overline{\sigma_i^2} + \delta_{i\sigma} \ln(\sigma_{i,t-1}^2) + \tau_{i\sigma} \eta_{i\sigma,t} \quad \eta_{i\sigma,t} \sim N(0, 1) \quad (6)$$

where $\bar{\cdot}$ indicates the unconditional mean, and δ_{ij} the persistence of the exogenous shocks $\eta_{ij,t}$, for $j = 0, \dots, K, \sigma$. This dynamics is sufficiently general to include a variety of existing specifications and capture a wide set of economic scenarios. Frequent shocks, such as changing firm characteristics, can be captured through quickly mean-reverting betas, while a high level of persistence δ_{ij} captures those shocks at low frequencies, (e.g. changes in the level of leverage).

2.2 Prior Specification

For parameter inference in (3)-(6) I opt for a Bayesian approach. Such an approach allows to characterize virtually any function of the model structural parameters (e.g. cross-sectional R^2 , variance ratios, long-run time series pricing errors, etc.). For each asset/portfolio the model parameters are $\theta_i = (\delta_{ij}, \overline{\beta_{ij}}, \tau_{ij}^2, \delta_{i\sigma}, \ln \overline{\sigma_i^2}, \tau_{i\sigma}^2)$. As the model is linear in both states and parameters I consider conjugate priors. For the parameters of the alphas and risk exposures

$\beta_{ij,t}$, I use a Normal-Inverse-Gamma prior structure (see West and Harrison 1997);

$$(\delta_{ij}, \overline{\beta_{ij}}, \tau_{ij}^2) \sim NIG \left(\underline{m}_{\beta}^{ij}, \underline{B}_{\beta}^{ij}, \underline{\nu}_{\beta}^{ij}/2, \underline{\nu}_{\beta}^{ij} \underline{s}_{\beta}^{ij}/2 \right) \quad i = 1, \dots, N \quad (7)$$

where $\underline{m}_{\beta}^{ij}, \underline{B}_{\beta}^{ij}$ represent the location and scale hyper-parameters of the normal distribution, $\underline{\nu}_{\beta}^{ij}$ the initial degrees of freedom and $\underline{\nu}_{\beta}^{ij} \underline{s}_{\beta}^{ij}$ the scale parameter of an inverse-gamma distribution.

I assume the same conjugate Normal-Inverse-Gamma prior for idiosyncratic risks;

$$(\delta_{i\sigma}, \ln \overline{\sigma_i^2}, \tau_{i\sigma}^2) \sim NIG \left(\underline{m}_{\sigma}^i, \underline{B}_{\sigma}^i, \underline{\nu}_{\sigma}^i/2, \underline{\nu}_{\sigma}^i \underline{s}_{\sigma}^i/2 \right) \quad i = 1, \dots, N \quad (8)$$

The cross-sectional pricing error and the set of risk premia $\gamma_t = (\gamma_{0,t}, \gamma'_{1:K,t})$ are not dynamic in nature. Indeed, time-variation is inherited from the instability of the factors risk exposures. Thus, the estimation of the no-arbitrage restriction can be simplified as a multi-variate linear regression, conditional on the betas. Given the independence structure of the error term $e_{i,t}$ in (4), I assume again a conjugate prior structure of the form

$$(\gamma, \omega^2) \sim NIG \left(\underline{\gamma}, \underline{\Gamma}, \underline{\omega}/2, \underline{\omega n}/2 \right) \quad (9)$$

with $\underline{\gamma}, \underline{\Gamma}, \underline{\omega}$ and $\underline{\omega n}$ the priors location and scale hyper-parameters.

2.3 Posterior Simulation

Posterior results are obtained through the Gibbs sampler algorithm developed in Geman and Geman (1984) in combination with the data augmentation technique by Tanner and Wong (1987) and Frühwirth-Schnatter (1994). The latent variables $B = \{\beta_{i,t}\}_{i=1,t=1}^{N \times T}$ and $\Sigma = \{\sigma_{i,t}^2\}_{i=1,t=1}^{N \times T}$, are simulated alongside the model parameters $\theta = \{\theta_i\}_{i=1}^N$ and the risk premia $\gamma = \{\gamma_{0,t}, \gamma'_{1:K,t}\}_{t=1}^T$. The initial step of the Gibbs sampler is to characterize the complete likelihood function, namely, the joint density of the data and the state variables.

$$p(Y, B, \Sigma | \theta, \gamma, F) = \prod_{t=1}^T \left(\prod_{i=1}^N p(y_{it} | F_t, \beta_{i,t}, \sigma_{i,t}^2, \gamma) p(\sigma_{i,t}^2 | \sigma_{i,t-1}^2, \delta_{i\sigma}, \overline{\sigma_i^2}, \tau_{i\sigma}^2) \right. \\ \left. \times \prod_{j=0}^K p(\beta_{ij,t} | \beta_{ij,t-1}, \delta_{ij}, \overline{\beta_{ij}}, \tau_{ij}^2) \right), \quad (10)$$

where $F = \{F_t\}_{t=1}^T$, and $Y = \{y_t\}_{t=1}^T$ with $y_t = (y_{1,t}, y_{2,t}, \dots, y_{N,t})$ the vector of assets/portfolios excess returns. Combining the priors with the complete likelihood, we obtain the posterior density $p(\theta, B, \Sigma, \gamma | R, F) \propto p(\theta, \gamma)p(R, B, \Sigma | \theta, \gamma, F)$. The Gibbs sampler is a combination of the Forward Filtering Backward Sampling (FFBS) of Carter and Kohn (1994) and Omori, Chib, Shepard, and Nakajima (2007). At each iteration of the sampler we sequentially cycle through the following steps:

1. Draw B conditional on Σ, θ, R and F .
2. Draw θ_β conditional on B, R and F .
3. Draw Σ conditional on B, θ, R and F .
4. Draw θ_σ conditional on B, R and F .
5. Draw γ conditional on B and θ .

I use a burn-in period of 2,000 and draw 10,000 observations storing every other of them to simulate the posterior distribution of parameters and latent variables. The resulting auto-correlations of the draws are very low.² A more detailed description of the Gibbs sampler and the corresponding convergence properties are given in Appendix D.

2.4 Testing Pricing Restrictions

Tests of conditional factor models boil down to answer questions like: “What is the probability that the pricing error is null at time t , given the in-sample information available up to time t ?”. Existing methods mostly rely on unconditional moments, biased-corrected asymptotic properties of estimators, and treat parameter estimates as the *true* values in the returns generating process. Given the sample size commonly used in testing factor pricing models, it remains unclear if such asymptotic inference is reliable in finite-samples. This is true even if we correct for error-in-measurement, heteroskedasticity, selection biases, etc. Yet, the assumption that parameters are observable is objectively restrictive. Indeed, in a classical setting, inference on pricing errors is to be read as contingent on the econometrician having full confidence in his parameters estimates, which is, of course, rarely the case.

²In order to gain a rough idea of how well the chain mixes in our algorithm we follow Primiceri (2005) and check the autocorrelation function of the draws.

In this paper, I develop an exact finite sample methodology to dynamically test conditional factor models fully acknowledging uncertainty on the structural parameters of the returns generating process. Such testing methodology can be applied both on single assets and jointly across portfolios. Yet, I develop the method to test conditional factor models both in the time series and in the cross-section. More specifically, I test the hypothesis that (3) holds at a particular point in time t , without requiring the restriction to be imposed at any other time. In the time series, I test the null that the pricing errors are not statistically significant, both on single portfolios in isolation, $\mathcal{H}_0 : \beta_{i0,t} \neq 0$ against the alternative $\mathcal{H}_1 : \beta_{i0,t} = 0$, and jointly across assets $\mathcal{H}_0 : \beta_{10,t} = \beta_{20,t} = \dots = \beta_{N0,t} = 0$. These can be written more compactly as $\mathcal{H}_0 : H\beta_t = q$, with H an $R \times N(K+1)$ selection matrix, q an R -dimensional vector of zeros, R the number of restrictions, $\beta_t = (\beta'_{0,t}, \beta'_{1,t}, \dots, \beta'_{K,t})$ and $\beta'_{j,t} = (\beta_{1j,t}, \beta_{2j,t}, \dots, \beta_{Nj,t})$ the risk exposures on the j th factor across assets/portfolios.

Hypothesis testing can be handled by computing standard Bayes factors. A convenient way to calculate Bayes factors comparing a restricted, \mathcal{H}_0 , to an unrestricted model, \mathcal{H}_1 , is the so-called Savage-Dickey density (SDD) ratio (see Verdinelli and Wasserman 1995). The SDD ratio represents the ratio of the posterior $p(\beta_{10,t} = \beta_{20,t} = \dots = \beta_{N0,t} = 0 | Z^t, \mathcal{H}_1)$ and prior $p(\beta_{10,t} = \beta_{20,t} = \dots = \beta_{N0,t} = 0 | \mathcal{H}_1)$ marginal probabilities. Posterior probabilities are computed condition on information available up to time t , $Z^t = (Y^t, F^t)$, with $Y^t = (y_1, \dots, y_t)$ and $F^t = (F_1, \dots, F_t)$ the recursive information about portfolios returns and risk factors, respectively. The main advantage of the SDD ratio is that involves only manipulation of priors and posteriors for the conditional factor pricing model. These are readily available from the estimation output. If the prior structure is common for both the model under the null \mathcal{H}_0 , and the alternative \mathcal{H}_1 , the following relationship is satisfied;

$$p(\beta_t | H\beta_t = q, \mathcal{H}_1) = p(\beta_t | \mathcal{H}_0)$$

then the Bayes factor comparing the null \mathcal{H}_0 and the alternative \mathcal{H}_1 can be computed as;

$$\mathcal{BF}_{0,1}^t = \frac{p(H\beta_t = q | Z^t, \mathcal{H}_1)}{p(H\beta_t = q | \mathcal{H}_1)} \quad (11)$$

Both the numerator and the denominator can be conveniently sampled from the output of the

MCMC scheme. The marginal posterior and prior probabilities in (11) are defined as

$$p(H\beta_t = q|Z^t, \mathcal{H}_1) = \int p(H\beta_t = q|Z^t, \mathcal{H}_1, \theta) p(\theta|Z^T) d\theta \quad (12)$$

$$p(H\beta_t = q|\mathcal{H}_1) = \int p(H\beta_t = q|\mathcal{H}_1, \theta) p(\theta|Z^T) d\theta \quad (13)$$

Conditional on idiosyncratic risks and the $N(K + 1)$ -dimensional vector of parameters θ , it can be shown that

$$p(H\beta_t = q|Z^t, \mathcal{H}_1, \theta) = N(Hm_t, HC_tH') \quad (14)$$

$$p(H\beta_t = q|\mathcal{H}_1, \theta) = N(H\underline{\beta}_t, H\underline{V}_tH') \quad (15)$$

where $m_t = E[\beta_t|Z^t, \theta]$, and $C_t = Var[\beta_t|Z^t, \theta]$ the mean and variance of factors risk exposures conditioned on the information available upto time t , while for the prior

$$\underline{\beta}_t = \bar{\beta} + \delta^t m_0 \quad (16)$$

$$\underline{V}_t = \delta^t C_0 (\delta^t)' + \sum_{j=0}^{t-1} \delta^j \tau (\delta^j \tau)' \quad (17)$$

for initial values m_0 and C_0 (see Appendix B for more details). Under the assumption of correct specification of conditional factor models, the intercept $\gamma_{0,t}$ should not be statistically different from zero. Therefore, a cross-sectional test involves to investigate the null $\mathcal{H}_0 : H\gamma_t = q$, against the alternative $\mathcal{H}_1 : H\gamma_t \neq q$, where H is a $R \times (K + 1)$ selection matrix, q an R -dimensional vector of zeros, R the number of restrictions (i.e. $R = 1$ in this case). Again, hypothesis testing can be handled by computing a Bayes factor comparing the unrestricted, \mathcal{H}_1 , to the restricted \mathcal{H}_0 cross-sectional regression at time t ;

$$\mathcal{BF}_{0,1}^t = \frac{p(H\gamma_t = q|\beta_t, \mathcal{H}_1)}{p(H\gamma_t = q|\mathcal{H}_1)} \quad (18)$$

Given the independence of γ_t estimates across time, both the numerator and the denominator can be conveniently sampled from the output of the MCMC scheme. The marginal posterior

and prior probabilities in (18) are defined as

$$p(H\gamma_t = q|\beta_t, \mathcal{H}_1) = \int p(H\gamma_t = q|\beta_t, \mathcal{H}_1, \theta) p(\theta|Z^T) d\theta \quad (19)$$

$$p(H\gamma_t = q|\mathcal{H}_1) = \int p(H\gamma_t = q|\mathcal{H}_1, \theta) p(\theta|Z^T) d\theta \quad (20)$$

where $\theta = \omega^2$. From the conjugate prior structure showed above it is easy to see that

$$p(H\gamma_t = q|\beta_t, \mathcal{H}_1, \theta) = N(H\bar{\gamma}, H(\omega^2\bar{\Gamma})H') \quad (21)$$

$$p(H\gamma_t = q|\mathcal{H}_1, \theta) = N(H\underline{\gamma}, H(\omega^2\underline{\Gamma})H') \quad (22)$$

where

$$\begin{aligned} \bar{\Gamma} &= (\underline{\Gamma}^{-1} + X'_\gamma X_\gamma)^{-1} \quad \text{with} \quad X_\gamma = [t, \beta_{1:N,t}] \\ \bar{\gamma} &= \bar{\Gamma} (\underline{\Gamma}^{-1}\underline{\gamma} + X'_\gamma y_t) \end{aligned}$$

with $y_t = (y_{1,t}, \dots, y_{N,t})$, and $\beta_{1:N,t}$ the $(N \times K)$ matrix of betas of each portfolio on each factor, except the Jensen's alpha $\beta_{i0,t}$. From the Bayes factors, (11) and (18), we can compute posterior probabilities of the null hypothesis $p(\mathcal{H}_0|Z^t)$. Assuming equal prior over the null and the alternative hypothesis, $p(\mathcal{H}_0) = p(\mathcal{H}_1)$, we can compute (see Robert 2007, Ch.5);

$$p(\mathcal{H}_0|Z^t) = \left[1 + \frac{1}{\mathcal{BF}_{0,1}^t} \right]^{-1} = \frac{\mathcal{BF}_{0,1}^t}{1 + \mathcal{BF}_{0,1}^t} \quad (23)$$

Note $p(\mathcal{H}_0|Z^t)$ might be interpreted as a p-value. Unlike standard p-value, however, the posterior probability naturally penalizes for the complexity of the model, being a direct function of the Bayes factor $\mathcal{BF}_{0,1}^t$. Same this applies for $p(\mathcal{H}_0|\beta_t)$. This addresses the so-called Lindleys paradox which is the apparent conflict between standard frequentist and Bayesian hypothesis testing. The conflict arises since standard t-statistics and corresponding p-values tend to go in favour of the null as the sample size increases.

3 Empirical Analysis

Following McCulloch and Rossi (1991), Geweke and Zhou (1996) and Bianchi et al. (2014), I jointly estimate pricing errors, risk exposures and premia in a single step, then overcome the issues of standard two-step methodologies. The system (3)-(6) is estimated via Bayesian single-step estimation scheme detailed in Appendix A. As application, I consider three main factor models common in the empirical finance literature. The first model is a conditional version of the Sharpe-Lintner CAPM. The CAPM performed well in early tests but has fared poorly since (see e.g. Fama and MacBeth 1973). The second model, which extends the CAPM by including two others risk factors, is the Fama and French (1993) three-factor model (FF3 henceforth) which includes two additional risk factors, beyond market risk, to proxy for size and value effects. The third model is the four-factor model (FF4 henceforth) proposed in Charhart (1997), which extends the three-factor model of Fama and French (1993) by including a portfolio to capture momentum effects on stock returns.

3.1 Data

The empirical analysis focuses on a set of standard test portfolios sorted by size, book-to-market ratio and momentum. I use the twenty-five portfolios constructed by double-sorting the stocks of NYSE/AMEX/NASDAQ along size and book-to-market dimensions. The portfolios are constructed at the end of each June using the corresponding market equity and NYSE breakpoints. Size is the market value of equity at the end of each June, and Book-to-Market (BM) is the ratio of book equity for the last fiscal year to market equity at December of the same year. In addition, I use deciles portfolios sorted on past performances. These momentum portfolios are constructed sorting every month into deciles stocks based on past six-month realized returns, as in Jegadeesh and Titman (1993). The portfolio returns are value-weighted averages in each group.³ The aggregate market portfolio is represented by the value-weighted NYSE/AMEX/NASDAQ index, taken from the Center for Research in Security Prices (CRSP). I use monthly returns in excess of the 1-month T-bill rate. The *SMB*, *HML* and *UMD* are taken from Kenneth French's website. *SMB* represents the return spread between portfolios of stocks

³The data are obtained from the Kenneth French's website

with small and large market capitalization, while *HML* is the return difference between “value” and “growth” stocks, namely portfolios of stocks with high and low book-to-market ratios. Finally, *UMD* is a zero-cost portfolio that is long previous 12-month return winners and short previous 12-month loser stocks. The sample period is 1963:07-2013:12. For the sake of exposition the deciles momentum portfolios are clustered in quintiles. The first ten years of monthly observations (1963-1973) are used to calibrate the hyper-parameters of the prior distributions. Given that these portfolios are widely used in the literature, I omit providing summary statistics for their returns. Table 1 reports usual OLS estimates of unconditional alphas across factor models. Standard errors are corrected for both heteroskedasticity and autocorrelation. Bold-faced numbers denote estimates statistically significant at the 5% confidence level. Panel A shows the pricing error for the CAPM.

[Insert Table 1 about here]

After adjusting for market risk, small stocks tend to show strong unexplained excess returns. Interestingly, large stocks do not show statistical significant unconditional pricing error. The alpha of past loser stocks change in sign turning negative. Panel B shows the same unconditional pricing errors according to the three-factor Fama-French model. As we would expect from previous literature, the amount of unexplained returns sensibly drops. However, both growth stocks and momentum-sorted portfolios do not seem to be sensibly priced. Third panel shows the results of the four-factor model FF4. Again, the amount unexplained returns sensibly reduces, albeit statistically significant for 6 out 30 portfolios.

3.2 Prior Choices and Parameters Estimates

Realistic values for the different prior distributions obviously depend on the problem at hand. The prior belief on δ_{ij} are such that changes in $\beta_{ij,t}$ s are highly persistent and the impact of unexpected news $\eta_{ij,t}$ can potentially be large. This view is consistent with both economic theory and previous empirical studies. Gomes, Kogan, and Zhang (2003) suggest that risk exposures are mainly a function of productivity shocks which take place at a business cycle frequency. Santos and Veronesi (2006) argue that stock betas are driven by the ratio of labor income

to total consumption, which is also highly persistent. Petkova and Zhang (2005) confirmed empirically that the market beta reverts to its long-term mean over a period of time consistent with the business cycle. Ang and Chen (2007) show that conditional betas are quite persistent once investors uncertainty is taken into account. Finally, Adrian and Franzoni (2009) suggest that low frequency fluctuations in the ex-ante market risk exposure can effectively support the conditional CAPM for value portfolios. Since portfolios returns are linearly dependent on the betas, to ensure stationarity the persistence parameter should be less than one in absolute value, $|\delta_{ij}| < 1$, for $j = 0, 1, \dots, K$. Note, (5) nests a model with perfect mean reversion (i.e. $\delta_{ij} = 0$). A non-parametric dynamics as in Campbell and Voulteenaho (2004), Fama and French (2005), and Lewellen and Nagel (2006), can then be seen as a special case of the model (3)-(6) with $\delta_{ij} = 0$ for $j = 0, 1, \dots, K$ and $i = 1, \dots, N$.

In order to reduce the sensitivity of posterior estimates to the prior specification, I use the initial ten years of monthly observations on portfolios and factor to train the hyper-parameters of weakly information priors. Table 2 reports the posterior estimates of the parameters from the three-factor Fama-French model. In order to keep the results readable, I report the results for only six portfolios: Small-Growth, Small-Value, Large-Growth, Large-Value, Loser and Winner. Top panel shows that the news $\eta_{i0,t}$ do not have a strong persistent effect on $\beta_{i0,t}$, meaning the Jensen's alphas shows a strong mean-reverting path across the sample. Interestingly, the unconditional alphas $\overline{\beta_{i0}}$ are statistically significant for just two out six portfolios. Despite their low duration, news on the time series pricing error have a strong short-term effect as the scale parameters τ_{i0} are large, spanning from 1.745 for Small-Value to 1.544 for the Small-Growth portfolio.

[Insert Table 2 about here]

Second panel shows that deviations of market risk exposures to their unconditional mean are much more persistent. The monthly autocorrelation of the market betas ranges from 0.93 for the Large-Growth portfolio, to 0.96 for the Small-Value stocks. The impact of the news τ_{i1} is sensibly lower, on average, than for the conditional alphas. The high persistence of the betas is consistent with both economic theory and previous empirical studies (e.g. Gomes et al. 2003, Petkova and Zhang 2005, Jostova and Philipov 2005, Santos and Veronesi 2006, Ang and Chen 2007, and Adrian and Franzoni 2009). Third and fourth panels show the results for conditional betas on

size and value risk factors, respectively. In general, both SMB and HML show highly persistent betas, with a marginally higher effect of unexpected news on both. Finally, bottom panel shows that idiosyncratic risks are highly persistent as well. The parameters estimates in Table 2, show that conditional alphas are likely affected by frequent and sizable events, such as short-lived changing economic conditions or firm’s characteristics. On the other hand, risk exposures and idiosyncratic risks are more likely to be affected by smaller and less transitory events, such as permanent changes in productivity, leverage or cash flows.

3.3 Pricing Errors

3.3.1 Jensen’s Alpha. A first direct evaluation on conditional factor models is to test the null that the joint alphas are not statistically different from zero conditioning on the information available upto time t . I test this restriction across both the CAPM, the three-factor Fama-French model and the FF4 specification. Figure 1 reports the posterior probability of the null \mathcal{H}_0 computed from (11) and (23), for each of the factor models and across the sample 1973:08-2013:01. Top panel shows the results for the CAPM. Except occasional nuances, the null hypothesis that the CAPM explain the time series of portfolios returns can not be sensibly rejected. The posterior probability of the null hypothesis $p(\mathcal{H}_0|Z^t)$ is indeed well above the standard 5% threshold value.

[Insert Figure 1 about here]

Mid panel shows the results from the Fama-French three-factor model. The FF3 safely shows absence of systematic mispricing across the time series of test portfolios. The results from the four-factor model FF4 reported in the bottom panel, confirm the performance of the CAPM and the FF3. The overall evidence is that multi-factor models provide a fairly accurate description of the in-sample time-series variation across the test portfolios. A further inspection of assets specific Jensen’s alphas confirm this conclusion. Figures 2-4 report the posterior probabilities of the no-pricing error null hypothesis for each of the portfolios under investigation. For the sake of exposition, I report the results for only six portfolios: Small-Growth, Small-Value,

Large-Growth, Large-Value, Loser and Winner. Figure 2 shows the results for the CAPM.

[Insert Figure 2 about here]

Except for occasional nuances, the null hypothesis that the CAPM explain the time series of portfolios returns cannot be systematically rejected. In other words, there are no sensible evidences of systematic deviations of test portfolios to mean-variance efficiency in a conditional sense. However, the performance of the CAPM tends to deteriorate across the recent financial crisis, and around 2004 for the Small-Value portfolio. Figure 3 and 4 show the results for the three- and four-factor model, respectively.

[Insert Figure 3 and 4 about here]

Both of the conditional factor models considered perform remarkably well. The null hypothesis of no-pricing error is practically never rejected at the 1% confidence level across the sample. As a whole, multi-factor models such as FF3 and FF4 do show explanatory power on the in-sample variation of test portfolios, at least in the time series. Table 3 confirms these results. Top panel shows the results for the CAPM. The average alphas across the test portfolios is never statistically different from zero. The in-sample variability of conditional estimates, as shown by the standard errors, is quite high. This means that, despite pricing error might be significant in terms of magnitude, its inherent uncertainty does not allow to claim any clearly detectable mispricing. The in-sample variation of conditional alphas is particularly important for the conditional CAPM, while for the conditional three- and four-factor models standard errors sensibly drop.

[Insert Table 3 about here]

Despite a lower variability, however, the uncertainty of the Jensen's alphas estimates is still way too high to sensibly detect (average) mispricing in the time series of the test portfolios. The four-factor model (bottom panel) draws the same conclusion of the FF3 model. Such higher volatility in the Jensen's alphas is consistent with previous empirical evidence such as Ang and Chen (2007), Ang and Kristensen (2012), and Bali and Engle (2014). Figures 2-4 suggest that conditional factor model in its specification (3)-(6) are consistent with mean-variance efficiency.

Although informative, these evidences are not exhaustive as other terms are involved in the observationally equivalent unconditional alpha (see Lewellen and Nagel (2006) and Lewellen et al. (2010)). As far as the dynamics of risk exposures is stationary, and under the null that the model is correct, the long-run pricing error coincides with the unconditional means of the alphas, $\overline{\beta_{i0}}$. Table 4 reports the estimation results.

[Insert Table 4 about here]

Top panel shows the results for the conditional CAPM. The unconditional pricing error turn out to be statistically significant on 6 out 30 test portfolios at the 5% confidence level. Past loser stocks as well as small-value stocks sensibly deviates from the unconditional mean-variance frontier. Mid and bottom panels show the model-implied unconditional alphas for both the FF3 and FF4, respectively. There is little evidence of significant unconditional pricing errors. Indeed, 2 out of 30 portfolios show a significant pricing error. Interestingly, there is no longer evidence of a value anomaly across the test portfolios.

3.4 Cross-Sectional Pricing Error

The $K + 1$ -dimensional vector of pricing error and risk premia $\gamma_t = (\gamma_{0,t}, \gamma'_{1:K,t})$ is not time-varying *per se*. Time variation is inherited from the stochastic nature of factor risks exposures. Indeed, each of the risk premia is computed directly from the no-arbitrage restriction (2). As an example, Figure 5 shows the no-arbitrage implied estimates of the conditional market risk premium.

[Insert Figure 5 about here]

Top panel reports the posterior median of the market risk premium $\gamma_{MKT,t}$. Clusters of instability can be found around late 80s, the period 2001/2002 and the recent great financial crisis. In fact, bottom panel shows that the conditional variance of the market risk premium spikes around these periods. Interestingly, Figure 5 shows that no clear predictability path arises from the no-arbitrage restriction. The model-implied conditional variance of the market risk premium is consistent with direct GARCH(1,1) estimates on the observable market excess returns, as shown by the red line on the bottom panel. Figure 6 reports the posterior probability of the

null $\mathcal{H}_0 : H\gamma_t = q$ computed from (18) and (23), across the sample 1973:08-2013:01. Top panel shows the posterior probability that the conditional CAPM holds in the cross-section at time t . The figure makes rather clear that the conditional CAPM does not explain the cross-sectional in-sample variation of the test portfolios. This is particularly true across the early 80s, the period 2001/2002 (i.e. 9/11, Financial scandals, Iraq war) and across the recent great financial crisis.

[Insert Figure 6 about here]

As far as the three-factor model is concerned (mid panel), the null hypothesis that $\gamma_{0,t} = 0$ is largely rejected. Except for occasional nuances, the “p-value” is indeed well above the 5% threshold. Remarkably, the performance of the FF3 model does not deteriorate during the recent great financial crisis. Finally, the bottom panel shows the results for a conditional specification of the four-factor FF4 model. Again, there is evidence of cross-sectional mispricing across the sample, especially during the period 2001/2002. As a whole, while the CAPM does not provide sensible results, the Fama-French model does not produce systematic and persistent deviation from mean-variance efficiency. In fact, apart from occasional fluctuations, separate calculations show that dynamic posterior probability imply the null is not rejected for more than 99% of the sample.

4 Sources of Instability and the Role of Idiosyncratic Risk

Lewellen and Nagel (2006) argue that the co-movement between the betas and both the conditional mean and variance of the market risk premium is not sufficient to justify unconditional pricing errors. In fact, although alphas and betas tend to vary over time, their instability would have to be implausibly large to explain unconditional asset pricing anomalies such as size, value and momentum premia. This inference is based on conventional OLS estimates, which rely on the assumption that market risk premium is stable enough and the market risk exposure is constant within sub-periods. However, both of these assumptions are violated. Figure 8 shows the median estimates of the conditional betas from the FF3 model, across six representative portfolios, and over the sample 1973-2013, monthly. Shaded areas represent the 95% credibility

regions.

[Insert Figure 7 about here]

Conditional betas can hardly be identified as stable over an annual/quarter horizon. Indeed, the median conditional beta for Large-Value stocks largely drops around the period 2001/2002. Losers and Winners portfolios tend to have quite a symmetric volatile exposure to aggregate market risk. For instance, conditional betas for the Loser portfolios show dramatic changes, climbing from 0.5 to 2 across the period 2001/2002 and the recent great financial crisis. These results seem to support the view expressed in Ang and Chen (2007). Such kind of instability is evident also for the SMB and HML betas. Figure 8 shows the results concerning the exposure to the SMB size factor.

[Insert Figure 8 about here]

As we would expected the exposure to SMB heavily fluctuates over time, and is not statistically significant for momentum-sorted portfolios. Interestingly, SMB turns out to have no effect also on Large-Value stocks, while largely affect both Growth and Small stocks. Similarly, momentum-sorted portfolios are not sensibly affected by the value risk proxy HML. In fact, the betas on the HML mimicking portfolio is not significant across the sample for both Loser and Winner stocks.

At the onset of the paper, I argue that conditional betas imperfectly reflects firm-specific news and broad economic conditions (e.g. Lettau and Ludvigson 2001). A simple way to check for such fundamental relation is to explore the correlation between the estimated betas and a set of well-established, observable, state variables. Table 5 studies the joint explanatory power of the standard predictors on the market betas. These state variables are the dividend yield (dy), the earnings-to-price ratio (ep), the dividend-payout ratio (dpr), net equity expansion (ntis), the default spread (def), log inflation (inf), industrial production (ip), the year-on-year consumption growth (cons), the M2 monetary aggregate (m2), the lagged excess return on the market ($ret_{(-1)}$), and the term spread (term). The sample period is 1973-2013, monthly. Data are from Goyal and Welch (2008) and the FREDII of the St. Louis Fed.

[Insert Table 5 about here]

Interestingly, the slopes on the (subset of) test portfolios show that conditional betas might reflect information on year-on-year consumption growth. In particular market risk exposure of small stocks tend to be counter-cyclical, while the opposite holds for large stocks. Past losers also show the well-established negative correlation with the business cycle. The impact of inflation and net equity issuance confirm the negative opposite effect of economic conditions on small vs. large stocks, as well as on past losers vs. past winners. Remarkably enough the adjusted R^2 span from 0.23 to 0.51, which means market betas effectively, although imperfectly, reflect news/information on economic conditions.

Above and beyond economic conditions, a second source of instability in the conditional betas comes from time variation in idiosyncratic risks. Intuitively, when portfolios returns are noisy, they become less informative on their exposure to aggregate market risk.⁴ As such, during high (low) volatility periods, recursive estimates of the betas tend to down-weight (over-weight) information from portfolios returns. Figure 7 shows the median estimates of the (square root of) idiosyncratic risks across six representative portfolios and over the sample 1973-2013, monthly. Shaded areas represent the 95% credibility regions.

[Insert Figure 7 about here]

As we would expect idiosyncratic risks spike around financial turmoils occurred across the period 2001/2002 and the recent great financial crisis. Surprisingly conditional variance is relatively flat across 2008/2009, while instead $\sigma_{i,t}^2$ climbs for value stocks and past winners. The relationship between the betas and idiosyncratic risks is particular evident for value and momentum portfolios. For instance, a spike in conditional variance coincide with a drop in the market risk exposure for large stocks and past winners. On the other hand, idiosyncratic risk is positively correlated with small stocks and past losers.

⁴The easiest way to understand the impact of time-varying idiosyncratic risk on the data informativeness is to consider the predictive likelihood

$$p(y_{i,t}|\beta_{i,t}, \sigma_{i,t}^2) = \frac{1}{\sqrt{2\pi\sigma_{i,t}^2}} \exp\left\{-\frac{1}{2\sigma_{i,t}^2} (y_{i,t} - \beta_{i0,t} - F_t' \beta_{i,t})^2\right\}$$

which is decreasing in the level of idiosyncratic risk σ_i^2 . In other words, when returns are noisy, the signal-to-noise ratio falls and data are less informative. This problem may vanish asymptotically but can be important in this setting due to small samples issues generated by the high persistence of $\beta_{i,t}$ and the relatively low signal-to-noise ratio.

5 Robustness Checks and Model Assessment

As a further model assessment of the in-sample explanatory power of different conditional factor models, I run both a variance decomposition test and a (log)marginal comparison across models. Under the hypothesis of a correct specification, conditional factor models should explain most of the predictable variation in the excess returns of the test portfolios, and leaving the unexplained portion as small as possible.⁵ I decompose at each time t the returns on test portfolios in a risk related $\gamma'_t\beta_{i,t}$ plus residual $\gamma_{0,t} + e_{i,t}$. From the output of the MCMC procedure I estimate the posterior distribution of the variance ratios as

$$\mathcal{VR}_{1,i} = \frac{\text{Var}[P(\gamma'_t\beta_{i,t}|\mathbf{Z}_{t-1})]}{\text{Var}[P(y_{i,t}|\mathbf{Z}_{t-1})]} > 0 \quad \mathcal{VR}_{2,i} = \frac{\text{Var}[P(\gamma_{0,t} + e_{i,t}|\mathbf{Z}_{t-1})]}{\text{Var}[P(y_{i,t}|\mathbf{Z}_{t-1})]} > 0. \quad (24)$$

for $i = 1, \dots, N$, where $P(\xi_t|\mathbf{Z}_{t-1})$ meaning *linear projection* of ξ_t onto the set of instrumental variables \mathbf{Z}_{t-1} . The $\mathcal{VR}_{1,i}$ should be equal to 1 if the conditional CAPM is correctly specified, which means that all the predictable variation in excess returns is captured by variation in market risk; at the same time, $\mathcal{VR}_{2,i}$ should be as close as possible to zero. I follow earlier literature such as Karolyi and Sanders (1998), Ferson and Harvey (1991) and Bianchi et al. (2014), and use a set of instrumental variables \mathbf{Z}_{t-1} used to tease out the total predictable variation in excess returns.⁶ Table 6 reports the results. For the sake of exposition, I report a representative subset of the test portfolios. Further results are available upon request.

[Insert Table 6 about here]

Columns 3 and 6 present posterior medians of $\mathcal{VR}_{1,i}$ and $\mathcal{VR}_{2,i}$ obtained from the conditional CAPM. The variance ratios are encouraging. On average, approximately 50% of the predictable variation in excess returns is captured by aggregate market risk. Explained predictable variation that however further increases for the three-factor Fama-French model (column 9). On average, the FF3 model specification explains around 80% of the in-sample predictable variation across

⁵Although (4) refers to excess returns, it remains a statistical approximation of the theoretical framework (2). This implies that in practice it may be naive to expect that market risk fully explain the cross-sectional variation of risk premia. A more sensible goal then seems to be that $\gamma'_t\beta_{i,t}$ explain the predictable variation in excess returns, rather the straight cross-sectional variance on itself.

⁶The set of instrumental variables consist of those detailed above. I also include a dummy variable to account for the so-called *January Effect* (see Thaler 1987).

portfolios. Column 15 reports the results for the four-factor model. Despite the increasing performance compared to the CAPM, the FF4 still underperforms the three-factor model. Indeed FF4, on average explains around 60% of the in-sample predictable variation across portfolios. Notably, $\mathcal{V}\mathcal{R}_{1,i}$ varies considerably across portfolios. The ratios are relatively higher large stocks and momentum portfolios.

Following McCulloch and Rossi (1991), I compute the marginal likelihood of the different factor model specifications to investigate their overall in-sample statistical performance. Indeed, the marginal likelihood of a model takes into account the latent nature of betas, idiosyncratic risks, as well as uncertainty on structural parameters. Intuitively, marginal likelihoods measure a model ability to explain the entire distribution of test portfolios. From the MCMC output the marginal likelihood of each model is computed as

$$p(R|F; \mathcal{M}_i) = \int \dots \int p(R|B, \Sigma, \theta; \mathcal{M}_i) p(\theta, B, \Sigma | R, F; \mathcal{M}_i) dB d\Sigma d\theta d\Sigma, \quad (25)$$

where \mathcal{M}_i identifies the i th model and the joint posterior density $p(\theta, B, \Sigma | R, F; \mathcal{M}_i)$ is computed through the Gibbs sampler (see Appendix A). As a direct function of marginal likelihoods Bayes' factors are used as model selection indicators that naturally penalizes the size/complexity of different models (e.g. Kass and Raftery 1995). Table 7 reports the (log) marginal likelihoods for different model specifications, as well as the (log) Bayes factors. Again, for the sake of exposition, I report a representative subset of the test portfolios. Further results are available upon request.

[Insert Table 7 about here]

The three-factor Fama-French model shows the higher (log) marginal likelihood, both as a whole (Global), and across the test portfolios. As pointed out in Kass and Raftery (1995), a log Bayes factor higher than 3/4 shows a decisive evidence in favor of model i versus j . Columns 5 to 7 shows that MKT, SMB and HML may effectively fit better the in-sample properties of the test portfolios. In fact, Bayes' factors are anywhere high than 4 in the comparison against the conditional CAPM, and, except for the *S4B5* portfolio, highly significant in comparison to the four-factor model FF4. In the light of these evidence, the ability to capture any predictable variation and the (log) marginal likelihood are directly related and raise the same conclusion.

6 Conclusion

I propose a methodology to jointly estimate time-varying factor loadings and premia in conditional factor models, and fully acknowledges parameter uncertainty in the portfolios returns generating process. I use Bayesian tools which helps generate posterior distributions of virtually any function of the structural parameters/statistics of the factor models. Joint estimates for time-varying betas, idiosyncratic risks and risk premia are implemented without having to instrument risk exposures and premia with observable predictors. I also provide a method to dynamically test conditional factor models both in the time series and in the cross-section. This allows to compute the posterior probability that the null of a factor model holds at time t , but does not necessarily hold at other times, and can be applied to both single assets and jointly across portfolios.

I apply the estimation and testing methodology to deciles portfolios sorted on size, book-to-market and past returns. In addition to the basic conditional CAPM, I test the well-known three-factor model of Fama and French (1993), together with the four-factor model proposed in Charhart (1997). I show that, once the dynamic and uncertain nature of the returns generating process is fully acknowledge, time series test do no reject the null/restriction of no pricing error across test portfolios. However, a dynamic cross-sectional tests show a superior performance of the three-factor Fama-French model. Indeed, conditional specifications of both the CAPM and the four-factor model reject the null of a sensible fit across the period 2001/2002 (i.e. 9/11, Financial scandals, Iraq war) and across the recent great financial crisis. Finally, a model comparison based on in-sample Bayes factor sensibly shows a better fit of the three-factor model compared to the CAPM and the four-factor model.

This work suggests some avenues for future research. Other asset pricing models not considered here could, of course, be examined. In terms of the methodology, the fact that we have to condition on latent stochastic volatility might represents a limitation which should be explored further. However, including such non-linearity may required the use of different technologies, such as, say, particle filters and I leave that for future developments. Finally, the estimation and testing framework can equivalently be applied incorporating other assets classes such as bonds and real estate (e.g. REITs).

References

- Adrian, T., and F. Franzoni. 2009. Learning About Beta: Time-Varying Factor Loadings, Expected Returns, and the Conditional CAPM. *Journal of Empirical Finance* pp. 537–556.
- Ang, A., and J. Chen. 2007. CAPM Over the Long-Run: 1926-2001. *Journal of Empirical Finance* pp. 1–40.
- Ang, A., and D. Kristensen. 2012. Testing Conditional Factor Models. *Journal of Financial Economics* pp. 132–152.
- Bali, T., and R. Engle. 2014. The Conditional CAPM Explains the Value Premium. *Working Paper* .
- Bianchi, D., M. Guidolin, and F. Ravazzolo. 2014. Macroeconomic Factors Strike Back: A Bayesian Change-Point Model of Time-Varying Risk Exposures and Premia in the U.S. Cross-Section. *Norges Bank Working Paper* .
- Bossaerts, P., and R. Green. 1989. A General Equilibrium Model of Changing Risk Premia: Theory and Tests. *Review of Financial Studies* 2:467–493.
- Campbell, J., and T. Voulteenaho. 2004. Bad Beta, Good Beta. *American Economic Review* 94:1249–1275.
- Campbell, S., and F. Diebold. 2009. Stock returns and expected business conditions: Half a century of direct evidence. *Journal of Business and Economic Statistics* pp. 266–278.
- Carter, C., and R. Kohn. 1994. On Gibbs sampling for state-space models. *Biometrika* pp. 541–553.
- Charhart, M. 1997. On Persistence in Mutual Fund Performance. *The Journal of Finance* 52:57–82.
- Cochrane, J. 2001. *Asset Pricing*. Princeton, NJ: Princeton University Press.
- Fama, E., and K. French. 1993. Common Risk Factors in the Returns on Stocks and Bonds. *Journal of Financial Economics* 33:3–56.
- Fama, E., and K. French. 1997. Industry Costs of Equity. *Journal of Financial Economics* pp. 153–193.
- Fama, E., and K. French. 2005. The Value Premium and the CAPM. *Working Paper, University of Chicago* .
- Fama, E., and J. MacBeth. 1973. Risk, Return and Equilibrium. *Journal of Political Economy* pp. 607–636.
- Ferson, W., and C. Harvey. 1991. The Variation of Economic Risk Premiums. *Journal of Political Economy* 99:385–415.
- French, K., G. Schwert, and R. Stambaugh. 1987. Expected Stock Returns and Volatility. *Journal of Financial Economics* pp. 3–29.
- Frühwirth-Schnatter, S. 1994. Data Augmentation and Dynamic Linear Models. *Journal of Time Series Analysis* 15:183–202.
- Geman, S., and D. Geman. 1984. Stochastic Relaxation, Gibbs Distributions, and the Bayesian Restoration of Images. *IEEE Transactions* 6:721–741.
- Geweke, J., and G. Zhou. 1996. Measuring the Pricing Error of the Arbitrage Pricing Theory. *Review of Financial Studies* 9:557–587.
- Gomes, J., L. Kogan, and L. Zhang. 2003. Equilibrium Cross-Section of Returns. *Journal of Political Economy* 4:693–732.
- Goyal, A., and I. Welch. 2008. A Comprehensive look at the empirical performance of equity premium prediction. *Review of Financial Studies* 4:1455–1508.
- Harvey, C. 2001. The Specification of Conditional Expectations. *Journal of Empirical Finance* pp. 573–638.
- Jagannathan, R., and Z. Wang. 1996. The Conditional CAPM and the Cross-Section of Expected Returns. *The Journal of Finance* 51:3–53.
- Jegadeesh, A., and S. Titman. 1993. Returns to Buying Winners and Selling Losers: Implications for Stock Market Efficiency. *The Journal of Finance* pp. 65–91.
- Jostova, G., and A. Philipov. 2005. Bayesian Analysis of Stochastic Betas. *Journal of Financial and Quantitative Analysis* 40:747–778.
- Kan, R., C. Robotti, and J. Shanken. 2013. Pricing Model Performance and the Two-Pass Cross-Sectional Regression Methodology. *The Journal of Finance* 6:2617–2649.
- Karolyi, G., and A. Sanders. 1998. The Variation of Economic Risk Premiums in Real Estate Returns. *Journal of Real Estate Finance and Economics* 17:245–262.

- Kass, R., and A. Raftery. 1995. Bayes Factors. *Journal of the American Statistical Association* 430:773–795.
- Koop, G., R. Leon-Gonzalez, and R. Strachan. 2010. Dynamic Probabilities of Restrictions in State Space Models: An Application to the Phillips Curve. *Journal of Business and Economic Statistics* 28:370–379.
- Lettau, M., and S. Ludvigson. 2001. Resurrecting the (C)CAPM: A Cross-Sectional Test When Risk Premia are Time Varying. *Journal of Political Economy* pp. 1238–1287.
- Lewellen, J., and S. Nagel. 2006. The Conditional CAPM Does Not Explain Asset-Pricing Anomalies. *Journal of Financial Economics* 82:289–314.
- Lewellen, J., S. Nagel, and J. Shanken. 2010. A Skeptical Appraisal of Asset Pricing Tests. *Journal of Financial Economics* 96:175–194.
- Liew, J., and M. Vassalou. 2000. Can book-to-market, size and momentum be risk factors that predict economic growth? *Journal of Financial Economics* pp. 221–245.
- McCulloch, R., and P. Rossi. 1991. A Bayesian Approach to Testing the Arbitrage Pricing Theory. *Journal of Econometrics* 49:141–168.
- Merton, R. 1973. An Intertemporal Capital Asset Pricing Model. *Econometrica* 41:867–887.
- Nardari, F., and J. Scruggs. 2007. Bayesian Analysis of Linear Factor Models with Latent Factors, Multivariate Stochastic Volatility, and APT Pricing Restrictions. *Journal of Financial and Quantitative Analysis* 42:857–891.
- Omori, Y., S. Chib, N. Shepard, and J. Nakajima. 2007. Stochastic Volatility with Leverage: Fast and Efficient Likelihood Inference. *Journal of Econometrics* 140:425–449.
- Pagan, A. 1984. Econometric Issues in the Analysis of Regressions with Generated Regressors. *International Economic Review* 25:221–247.
- Petersen, M. 2009. Estimating Standard Errors in Finance Panel Data Sets: Comparing Approaches. *Review of Financial Studies* 22:435–480.
- Petkova, R., and L. Zhang. 2005. Is Value Riskier Than Growth? *Journal of Financial Economics* 78:187–202.
- Primiceri, G. 2005. Time Varying Structural Vector Autoregressions and Monetary Policy. *Review of Economic Studies* 72:821–852.
- Robert, C. 2007. *The Bayesian Choice, Second Edition*. Springer.
- Santos, T., and P. Veronesi. 2006. Labour Income and Predictable Stock Returns. *The Review of Financial Studies* 19:1–44.
- Shanken, J. 1990. Intertemporal Asset Pricing: An Empirical Investigation. *The Journal of Econometrics* 45:99–120.
- Tanner, A., and W. Wong. 1987. The Calculation of Posterior Distributions by Data Augmentation. *Journal of the American Statistical Association* 82:528–540.
- Thaler, R. 1987. Anomalies: The January Effect. *Journal of Economic Perspectives* 1:197–201.
- Vassalou, M. 2003. News Related to future GDP growth as a risk factor in equity returns. *Journal of Financial Economics* pp. 47–73.
- Verdinelli, I., and L. Wasserman. 1995. Computing Bayes Factors Using a Generalization of the Savage-Dickey Density Ratio. *Journal of the American Statistical Association* pp. 614–618.
- West, M., and J. Harrison. 1997. *Bayesian forecasting and dynamics models*. Springer.

Appendix

A Gibbs Sampler

A.1 Step 1. Sampling the Conditional Alphas and Betas

The full conditional posterior density for the time-varying factor loadings is computed using a Forward Filtering Backward Sampling (FFBS) approach as in Carter and Kohn (1994). The initial prior are sequentially updated via the Kalman filtering recursion, and the parameters are drawn from the posterior distribution which is generated by backward recursion (see Frühwirth-Schnatter 1994, Carter and Kohn 1994, and West and Harrison 1997).

A.2 Step 2. Sampling the Parameters $\theta_\beta^{ij} = (\delta^{ij}, \overline{\beta}_{ij}, \tau_{ij}^2)$

I consider a Normal-Inverse-Gamma prior structure (see West and Harrison 1997).

$$(\delta_{ij}, \overline{\beta}_{ij}, \tau_{ij}^2) \sim \text{NIG} \left(\underline{m}_\beta^{ij}, \underline{B}_\beta^{ij}, \underline{\nu}_\beta^{ij}/2, \underline{\nu}_\beta^{ij} \underline{s}_\beta^{ij}/2 \right) \quad i = 1, \dots, N \quad (\text{A.26})$$

where $\underline{m}_\beta^{ij}, \underline{B}_\beta^{ij}$ the location and scale hyper-parameters of the normal distribution, $\underline{\nu}_\beta^{ij}$ the initial degrees of freedom and $\underline{\nu}_\beta^{ij} \underline{s}_\beta^{ij}$ the scale parameter of an inverse-gamma distribution. Posterior estimates are obtained once the factor loadings $\beta_{i,t}$ are sampled for each $t = 1, \dots, T$. Given the conjugate prior structure the updating scheme is easily derived as

$$(\delta^{ij}, \overline{\beta}_{ij} | \tau_{ij}^2, \beta_{i,1:T}) \sim N \left(\overline{m}_\beta^{ij}, \tau_{ij}^2 \overline{B}_\beta^{ij} \right) \quad (\text{A.27})$$

$$(\tau_{ij}^2 | \beta_{i,1:T}) \sim \text{IG} \left(\overline{\nu}_\beta^{ij}/2, \overline{\nu}_\beta^{ij} \overline{s}_\beta^{ij}/2 \right) \quad (\text{A.28})$$

for $i = 1, \dots, N$, with

$$\begin{aligned} \overline{B}_\beta^{ij} &= \left(\left(\underline{B}_\beta^{ij} \right)^{-1} + X_\beta' X_\beta \right)^{-1} \\ \overline{m}_\beta^{ij} &= \overline{B}_\beta^{ij} \left(\left(\underline{B}_\beta^{ij} \right)^{-1} \underline{m}_\beta^{ij} + X_\beta' Y_\beta \right) \\ \overline{\nu}_\beta^{ij} &= \underline{\nu}_\beta^{ij} + T \\ \overline{\nu}_\beta^{ij} \overline{s}_\beta^{ij} &= \underline{\nu}_\beta^{ij} \underline{s}_\beta^{ij} + \left(\underline{m}_\beta^{ij} \right)' \left(\underline{B}_\beta^{ij} \right)^{-1} \underline{m}_\beta^{ij} + Y_\beta' Y_\beta - \left(\overline{m}_\beta^{ij} \right)' \left(\overline{B}_\beta^{ij} \right)^{-1} \overline{m}_\beta^{ij} \end{aligned}$$

where $\beta_{i,1:T} = (\beta_{i,1}, \dots, \beta_{i,T})$, $X_\beta = [\iota, \beta_{i,1:T-1}]$, $Y_\beta = \beta_{i,2:T}$ and $\beta_{i,t} = (\beta_{i0,t}, \beta_{i1,t}, \dots, \beta_{iK,t})$.

A.3 Step 3 and 4. Sampling the Idiosyncratic Risk and the Corresponding Structural Parameters.

The conditional variances $\ln \sigma_{it}^2$ preserves the standard properties of state space models. From (??) the log of squared residuals for the i th asset can be defined as

$$\ln (y_{i,t} - \beta_{i0,t} - F_t' \beta_{i,t})^2 = \ln \sigma_{i,t}^2 + u_t \quad (\text{A.29})$$

where $u_t = \ln \varepsilon_t^2$ has a $\ln \chi^2(1)$. As in Omori et al. (2007), I approximate the $\ln \chi^2(1)$ distribution with a finite mixture of ten normal distributions, such that the density of u_t is given by

$$p(u_t) = \sum_{l=1}^{10} \varphi_l \frac{1}{\sqrt{\omega_l^2 2\pi}} \exp \left(-\frac{(u_t - \mu_l)^2}{2\omega_l^2} \right) \quad (\text{A.30})$$

with $\sum_{l=1}^{10} \varphi_l = 1$. The appropriate values for μ_l, φ_l and ϖ_l^2 can be found in Omori et al. (2007). At each step of the algorithm I simulate a component of the mixture at each time t . Given the mixture component I apply a Kalman filter method, such that idiosyncratic risk can be sampled as the betas. The initial prior of the log idiosyncratic volatility $\ln \sigma_{i,0}^2$ is normal with mean -1 and conditional variance equal to 10. To sample $\theta_\sigma^i = (\delta^{i\sigma}, \ln \bar{\sigma}_i^2, \tau_{i\sigma}^2)$ from its joint posterior $p(\delta^{i\sigma}, \ln \bar{\sigma}_i^2, \tau_{i\sigma}^2 | \ln \Sigma_i)$ I proceed as in Step 2, with $\ln \Sigma_i = \{\ln \sigma_{it}^2\}_{t=1}^T$. The same updating scheme is adopted by defining $X_\sigma = [\iota, \ln(\sigma_{1:T-1}^2)]$ and $Y_\sigma = \ln(\sigma_{2:T}^2)$.

A.4 Step 5. Sampling the Factors Risk Premia

Conditional on the risk exposures, the estimate of the risk premia coincide with a multivariate linear model with uncorrelated errors. I assume a conjugate prior structure of the form

$$(\gamma, \omega^2) \sim NIG(\underline{\gamma}, \underline{\Gamma}, \underline{\omega}/2, \underline{\omega n}/2) \quad (\text{A.31})$$

with $\underline{\gamma}, \underline{\Gamma}, \underline{\omega}$ and $\underline{\omega n}$ the priors location and scale hyper-parameters. Posterior estimates are then obtained by updating the prior structure as

$$(\gamma | \omega^2, \beta_t) \sim N(\bar{\gamma}, \gamma^2 \bar{\Gamma}) \quad (\text{A.32})$$

$$(\gamma^2 | \beta_t) \sim IG(\bar{\gamma}/2, \bar{\gamma n}/2) \quad (\text{A.33})$$

for $t = 1, \dots, T-1$, with

$$\begin{aligned} \bar{\Gamma} &= (\underline{\Gamma}^{-1} + X_\gamma' X_\gamma)^{-1} \\ \bar{\gamma} &= \bar{\Gamma} (\underline{\Gamma}^{-1} \underline{\gamma} + X_\gamma' Y_\gamma) \\ \bar{\omega} &= \underline{\omega} + N \\ \bar{\omega n} &= \underline{\omega n} + \underline{\gamma}' \underline{\Gamma}^{-1} \underline{\gamma} + Y_\gamma' Y_\gamma - \bar{\gamma}' \bar{\Gamma}^{-1} \bar{\gamma} \end{aligned}$$

where $\beta_t = (\beta_{1,t}, \dots, \beta_{N,t})$ represents conditional estimates of the factors risks exposures at time t for each of the N portfolios, $X_\gamma = [\iota, \beta_t]$ and $Y_\gamma = y_t$.

B Dynamic Testing Methodology

Testing the conditional CAPM boils down to test a joint restriction of the form

$$\mathcal{H}_0 : H\beta_t = q$$

$$\mathcal{H}_1 : H\beta_t \neq q$$

where H is a $R \times N(K+1)$ selection matrix, q an R -dimensional vector of zeros, R the number of restrictions, $\beta_t = (\beta_{0,t}, \beta_{1,t}, \dots, \beta_{K,t})$ and $\beta_{j,t}' = (\beta_{1j,t}, \beta_{2j,t}, \dots, \beta_{Nj,t})$ the risk exposures on the j th factor across assets/portfolios. Hypothesis testing can be handled by computing standard Bayes factors. A convenient way to calculate Bayes factors comparing a restricted, \mathcal{H}_0 , to an unrestricted model, \mathcal{H}_1 , is the so-called Savage-Dickey density (SDD) ratio (see Verdinelli and Wasserman 1995). The main advantage of the SDD ratio is that involves only manipulation of priors and posteriors for the conditional factor pricing model. These are readily available from the estimation output. If the prior structure is common for both the model under the null \mathcal{H}_0 , and the alternative \mathcal{H}_1 , the following relationship is satisfied;

$$p(\beta_t | H\beta_t = q, \mathcal{H}_1) = p(\beta_t | \mathcal{H}_0)$$

also the (marginal) likelihoods of the observed data are observationally equivalent (see Verdinelli and Wasserman 1995 for a more detailed discussion). Nested models allows to express the Bayes factor corresponding to

test the restriction as a ratio of ordinates,

$$\mathcal{BF}_{0,1}^t = \frac{p(H\beta_t = q|Z^t, \mathcal{H}_1)}{p(H\beta_t = q|\mathcal{H}_1)} \quad (\text{A.34})$$

where $Z^t = (X^t, Y^t)$ and $X^t = (X_1, \dots, X_t)$, $Y^t = (Y_1, \dots, Y_t)$ the sample information on risk factors and portfolios returns, respectively. Following Koop et al. (2010) the prior and posterior in (A.34) can be sampled from recursive filtering estimates of conditional moments $E(\beta_t|Z^t, \theta)$ and $Var(\beta_t|Z^t, \theta)$ and the output of the MCMC estimation scheme. Indeed, conditional on the idiosyncratic risks and risk premia, the state space model in (3)-(6) can be written by stacking single equations as

$$y_t = X_t\beta_t + \epsilon_t \quad \epsilon_t \sim N(0, \Sigma_t) \quad (\text{A.35})$$

$$\beta_t = (1 - \delta)\bar{\beta} + \delta\beta_{t-1} + \tau\eta_t \quad \eta_t \sim N(0, I_{p \times 1}) \quad (\text{A.36})$$

with $p = N(K + 1)$, $X_t = [1 \ F_t'] \otimes I_N$, $\beta_t = (\beta'_{0,t}, \beta'_{1,t}, \dots, \beta'_{K,t})$ and $\beta'_{j,t} = (\beta_{1j,t}, \beta_{2j,t}, \dots, \beta_{Nj,t})$ the risk exposures on the j th factor across assets/portfolios. Given the independence of risk exposures across factors and portfolios, the intercept $(1 - \delta)\bar{\beta}$ is an $N(K + 1)$ -dimensional vector with the ij element equal to $(1 - \delta_{ij})\bar{\beta}_{ij}$. Likewise, $\delta = \text{diag}(\delta_{10}, \dots, \delta_{N0}, \delta_{11}, \dots, \delta_{N1}, \dots, \delta_{1K}, \dots, \delta_{NK})$, and $\tau = \text{diag}(\tau_{10}, \dots, \tau_{N0}, \tau_{11}, \dots, \tau_{N1}, \dots, \tau_{1K}, \dots, \tau_{NK})$. Conditionally on idiosyncratic risk Σ_t and θ , and assuming an initial distribution $\beta_0|y_0 \sim N(m_0, C_0)$, it is straightforward to show that the (see West and Harrison 1997 for more details)

$$\begin{aligned} \beta_t|Y^{t-1}, X^{t-1}, \theta &\sim N(a_t, R_t) && \text{Propagation Density} \\ Y_t|Y^{t-1}, X^{t-1}, \theta &\sim N(f_t, Q_t) && \text{Predictive Density} \\ \beta_t|Y^t, X^t, \theta &\sim N(m_t, C_t) && \text{Filtering Density} \end{aligned}$$

with

$$\begin{aligned} a_t &= (1 - \delta)\bar{\beta} + \delta m_{t-1} && R_t = \delta C_{t-1} \delta' + \tau \\ f_t &= a_t' X_t && Q_t = X_t' R_t X_t + \Sigma_t \\ m_t &= a_t + K_t e_t && C_t = R_t - K_t Q_t K_t' \end{aligned} \quad (\text{A.37})$$

and $K_t = R_t X_t Q_t^{-1}$ and $e_t = y_t - f_t$, the so-called *Kalman gain* and the investors' forecasting error, respectively. From (A.37) we can see that the betas of the unrestricted model \mathcal{H}_1 are distributed as a multivariate Normal distribution

$$p(\beta_t|Z^t, \mathcal{H}_1, \theta) = (2\pi)^{-N} |C_t|^{-1/2} \exp\left\{-\frac{1}{2}(\beta_t - m_t)' C_t^{-1} (\beta_t - m_t)\right\}$$

Using standard results on multivariate Gaussian distributions

$$p(H\beta_t|Z^t, \mathcal{H}_1, \theta) = (2\pi)^{-N} |HC_t H'|^{-1/2} \exp\left\{-\frac{1}{2}(H\beta_t - Hm_{i,t})' (HC_t H')^{-1} (H\beta_t - Hm_t)\right\}$$

therefore, the model restriction $\mathcal{H}_0 : H\beta_t = 0$ is distributed as

$$p(H\beta_t = q|Z^t, \mathcal{H}_1, \theta) = (2\pi)^{-N} |HC_t H'|^{-1/2} \exp\left\{-\frac{1}{2}(q - Hm_{i,t})' (HC_t H')^{-1} (q - Hm_t)\right\}$$

Now given the output from the MCMC algorithm, the numerator in (A.34) can be numerically approximated as

$$\hat{p}(H\beta_t = q|Z^t, \mathcal{H}_1) = \frac{1}{G} \sum_{g=1}^G p(H\beta_t = q|Z^t, \mathcal{H}_1, \theta^{(g)})$$

with g the number of draws from the MCMC algorithm.⁷ The same strategy can be applied to evaluate the prior at the denominator in (A.34). The hierarchical prior can be derived from the initial distribution $p(\beta_0|Y_0, Z_0) = N(m_0, C_0)$ and the recursion of the state equation β_t in (A.36). As such,

$$\begin{aligned}\beta_t &= \bar{\beta} + \delta^t \beta_0 + \sum_{j=0}^{t-1} \delta^j \tau^{1/2} \eta_{t-j} \\ \beta_t &= \bar{\beta} + \delta^t m_0 + \delta^t C_0^{1/2} z + \sum_{j=0}^{t-1} \delta^j \tau^{1/2} \eta_{t-j}\end{aligned}\tag{A.38}$$

with $z \sim N(0, I_p)$ with $p = N(K+1)$. The recursive prior has a location parameter equal to

$$\underline{\beta}_t = E(\beta_t|\mathcal{H}_1, \theta) = \bar{\beta} + \delta^t m_0$$

and prior variance

$$\underline{V}_t = Var(\beta_t|\mathcal{H}_1, \theta) = \delta^t C_0 (\delta^t)' + \sum_{j=0}^{t-1} \delta^j \tau (\delta^j \tau)'$$

Therefore

$$p(\beta_t|\mathcal{H}_1, \theta) = (2\pi)^{-N} |\underline{V}_t|^{-1/2} \exp\left\{-\frac{1}{2}(\beta_t - \underline{\beta}_t)' \underline{V}_t^{-1} (\beta_t - \underline{\beta}_t)\right\}$$

such that

$$p(H\beta_t = q|\mathcal{H}_1, \theta) = (2\pi)^{-N} |H\underline{V}_t H'|^{-1/2} \exp\left\{-\frac{1}{2}(q - \underline{\beta}_t)' (H\underline{V}_t H')^{-1} (q - \underline{\beta}_t)\right\}$$

Again, given the output of the MCMC algorithm, I can approximate the marginal prior as

$$\hat{p}(H\beta_t = q|\mathcal{H}_1) = \frac{1}{G} \sum_{g=1}^G p(H\beta_t = q|\mathcal{H}_1, \theta^{(g)})$$

Under the assumption of correct specification of conditional factor models, the cross-sectional intercept $\gamma_{0,t}$ should not be statistically different from zero. Therefore, a cross-sectional test involves to investigate the null $\mathcal{H}_0 : H\gamma_t = q$, against the alternative $\mathcal{H}_1 : H\gamma_t \neq q$, where H is a $R \times (K+1)$ selection matrix, q an R -dimensional vector of zeros, R the number of restrictions (i.e. $R = 1$ in this case). As above, hypothesis testing can be handled by computing a Bayes factor comparing the unrestricted, \mathcal{H}_1 , to the restricted \mathcal{H}_0 cross-sectional regression at time t ;

$$\mathcal{BF}_{0,1}^t = \frac{p(H\gamma_t = q|\beta_t, \mathcal{H}_1)}{p(H\gamma_t = q|\mathcal{H}_1)}\tag{A.39}$$

Given the independence of γ_t estimates across time, both the numerator and the denominator can be conveniently sampled from the output of the MCMC scheme. Conditional on the parameters $\theta = \omega^2$ and the prior hyperparameters $\underline{\Gamma}$ and $\underline{\gamma}$, the posterior probability of the factors risk premia is equal to

$$p(\gamma_t|\beta_t, \mathcal{H}_1, \theta) = (2\pi)^{-(K+1)} |\omega^2 \bar{\Gamma}|^{-1/2} \exp\left\{-\frac{1}{2}(\gamma_t - \bar{\gamma})' (\omega^2 \bar{\Gamma})^{-1} (\gamma_t - \bar{\gamma})\right\}$$

where the hyperparameters $\bar{\Gamma}$ and $\bar{\gamma}$ are defined as

$$\begin{aligned}\bar{\Gamma} &= (\underline{\Gamma}^{-1} + X_\gamma' X_\gamma)^{-1} & \text{with} & & X_\gamma &= [I, \beta_{[-0],t}] \\ \bar{\gamma} &= \bar{\Gamma} (\underline{\Gamma}^{-1} \underline{\gamma} + X_\gamma' y_t)\end{aligned}$$

⁷As usual the number of draws can be chosen to optimize the accuracy of the approximation, and standard diagnostics can be used to analyze convergence properties.

with $y_t = (y_{1,t}, \dots, y_{N,t})$, and $\beta_{[-0],t}$ the $(N \times K)$ matrix of betas of each portfolio on each factor, except the Jensen's alpha $\beta_{i0,t}$. From the properties of the multivariate normal distribution it can be shown that

$$p(H\gamma_t | \beta_t, \mathcal{H}_1, \theta) = (2\pi)^{-(K+1)} \left| \omega^2 H \bar{\Gamma} H' \right|^{-1/2} \exp \left\{ -\frac{1}{2} (H\gamma_t - H\bar{\gamma})' (\omega^2 H \bar{\Gamma} H')^{-1} (H\gamma_t - H\bar{\gamma}) \right\}$$

such that the pricing restriction is distributed as

$$p(H\gamma_t = q | \beta_t, \mathcal{H}_1, \theta) = (2\pi)^{-(K+1)} \left| \omega^2 H \bar{\Gamma} H' \right|^{-1/2} \exp \left\{ -\frac{1}{2} (q - H\bar{\gamma})' (\omega^2 H \bar{\Gamma} H')^{-1} (q - H\bar{\gamma}) \right\}$$

Given the output of the MCMC algorithm, I can approximate the marginal prior as

$$\hat{p}(H\gamma_t = q | \beta_t, \mathcal{H}_1) = \frac{1}{G} \sum_{g=1}^G p(H\gamma_t = q | \beta_t, \mathcal{H}_1, \theta^{(g)})$$

The conditional prior $p(H\beta_t = q | \mathcal{H}_1, \theta)$ can be directly sampled from a multivariate normal distribution with prior hyperparameters $\underline{\Gamma}$ and $\underline{\gamma}$. Then its marginal $\hat{p}(H\gamma_t = q | \mathcal{H}_1)$ can be approximated as above from the output of the MCMC scheme. Assuming equal prior over the null and the alternative hypothesis, $p(\mathcal{H}_0) = p(\mathcal{H}_1)$, we can compute (see Robert 2007, Ch.5);

$$p(\mathcal{H}_0 | Z^t) = \left[1 + \frac{1}{\mathcal{BF}_{0,1}^t} \right]^{-1} = \frac{\mathcal{BF}_{0,1}^t}{1 + \mathcal{BF}_{0,1}^t} \quad (\text{A.40})$$

Note $p(\mathcal{H}_0 | Z^t)$ might be interpreted as a p-value. Unlike standard p-value, however, the posterior probability naturally penalizes for the complexity of the model, being a direct function of the Bayes factor $\mathcal{BF}_{0,1}^t$. This address the so-called Lindleys paradox which is the apparent conflict between standard frequentist and Bayesian hypothesis testing. The conflict arises since standard t-statistics and corresponding p-values tend to go in favour of the null as the sample size increases. The posterior probability (A.40) makes clear this is not the case under the methodology I propose.

C Variance Decomposition Tests

We use the posterior densities of the time series of factor loadings and risk premia to perform a number of tests that allow us to assess whether a posited asset pricing framework may explain an adequate percentage of excess asset returns. (4) decomposes excess asset returns in a component related to risk, represented by the term $\gamma'_t \beta_{i,t}$ plus a residual $\gamma_{0,t} + e_{i,t}$. In principle, a multi-factor model is as good as the implied percentage of total variation in excess returns explained by its first component, $\gamma'_t \beta_{i,t}$. However, here we should recall that even though (4) refers to excess returns, it remains a statistical implementation of the framework in (3). This implies that in practice it may be naive to expect that $\gamma'_t \beta_{i,t}$ be able to explain much of the variability in excess returns. A more sensible goal seems to be that $\gamma'_t \beta_{i,t}$ ought to at least explain the *predictable* variation in excess returns. We therefore follow earlier literature, such as Karolyi and Sanders (1998), and adopt the following approach. First, the excess return on each asset is regressed onto a set of M instrumental variables that proxy for available information at time $t-1$, \mathbf{Z}_{t-1} ,

$$y_{i,t} = \lambda_{i0} + \sum_{m=1}^M \lambda_{im} Z_{m,t-1} + \xi_{i,t}, \quad (\text{A.41})$$

to compute the sample variance of fitted values,

$$\text{Var}[P(y_{i,t} | \mathbf{Z}_{t-1})] \equiv \text{Var} \left[\lambda_{i0} + \sum_{m=1}^M \lambda_{im} Z_{m,t-1} \right], \quad (\text{A.42})$$

where the notation $P(y_{i,t} | \mathbf{Z}_{t-1})$ means “linear projection” of x_{it} on a set of instruments, \mathbf{Z}_{t-1} . Second, for each asset $i = 1, \dots, N$, a time series of fitted (posterior) risk compensations, $\gamma'_t \beta_{i,t}$, is regressed onto the instrumental

variables,

$$\gamma'_t \beta_{i,t} = \lambda'_{i0} + \sum_{m=1}^M \lambda'_{im} Z_{m,t-1} + \xi'_{i,t} \quad (\text{A.43})$$

to compute the sample variance of fitted risk compensations:

$$\text{Var} [P(\gamma'_t \beta_{i,t} | \mathbf{Z}_{t-1})] \equiv \text{Var} \left[\lambda'_{i0} + \sum_{m=1}^M \lambda'_{im} Z_{m,t-1} \right]. \quad (\text{A.44})$$

The predictable component of excess returns in (A.41) not captured by the model is then the sample variance of the fitted values from the regression of the residuals $\gamma_{0,t} + e_{i,t}$ on the instruments:

$$\text{Var} [\gamma_{0,t} + e_{i,t}] = \text{Var} [P(\lambda_{0,t} + e_{i,t} | \mathbf{Z}_{t-1})]. \quad (\text{A.45})$$

At this point, it is informative to compute and report two variance ratios, commonly called $VR1$ and $VR2$, after Ferson and Harvey (1991):

$$\mathcal{VR1} \equiv \frac{\text{Var} [P(\gamma'_t \beta_{i,t} | \mathbf{Z}_{t-1})]}{\text{Var} [P(y_{i,t} | \mathbf{Z}_{t-1})]} > 0 \quad (\text{A.46})$$

$$\mathcal{VR2} \equiv \frac{\text{Var} [P(\gamma_{0,t} + e_{i,t} | \mathbf{Z}_{t-1})]}{\text{Var} [P(y_{i,t} | \mathbf{Z}_{t-1})]} > 0. \quad (\text{A.47})$$

$VR1$ should be equal to 1 if the multi-factor model is correctly specified, which means that all the predictable variation in excess returns is captured by variation in risk compensations; at the same time, $VR2$ should be equal to zero if the multi-factor model is correctly specified. Importantly, when these decomposition tests are implemented using the estimation outputs obtained from the MCMC scheme, drawing from the joint posterior densities of the factor loadings $\beta_{i,t}$ and the implied risk premia γ_t , $i = 1, \dots, N$ and $t = 1, \dots, T$, and holding the instruments fixed over time, it is possible to compute $\mathcal{VR1}$ and $\mathcal{VR2}$ in correspondence to each of such draws and hence obtain their posterior distributions.⁸

⁸Notice that $\mathcal{VR1} = 1$ does not imply that $\mathcal{VR2} = 0$ and viceversa, because

$$\text{Var} [P(y_{i,t} | \mathbf{Z}_{t-1})] \neq \text{Var} [P(\gamma'_t \beta_{i,t} | \mathbf{Z}_{t-1})] + \text{Var} [P(\gamma_{0,t} + e_{i,t} | \mathbf{Z}_{t-1})].$$

Table 1. Unconditional Alphas

The table reports the unconditional alphas computed from different factor models on monthly excess returns of size, value, and momentum portfolios. Deciles momentum portfolios are clustered in quintiles for the sake of exposition. Bold-faced numbers denote estimates statistically significant at the 5% confidence level. The standard errors are corrected for autocorrelation and heteroschedasticity. The sample period is 1963:01-2013:12. The first ten years are cut to be consistent with the training sample used for the estimation of the conditional CAPM in section 3.

	Mean					St. Error				
Panel A: CAPM										
Size/BM	1	2	3	4	5	1	2	3	4	5
1	-0.488	0.258	0.375	0.560	0.635	0.237	0.193	0.183	0.188	0.205
2	-0.177	0.190	0.419	0.471	0.461	0.164	0.149	0.153	0.163	0.193
3	-0.134	0.267	0.334	0.389	0.637	0.128	0.123	0.139	0.155	0.187
4	0.013	0.122	0.225	0.338	0.362	0.125	0.121	0.149	0.147	0.167
5	-0.094	0.126	0.077	0.151	0.215	0.091	0.080	0.100	0.131	0.165
Mom	-0.578	0.020	0.000	0.179	0.324	0.173	0.101	0.071	0.069	0.105
Panel B: FF3										
Size/BM	1	2	3	4	5	1	2	3	4	5
1	-0.552	0.030	0.063	0.170	0.112	0.125	0.085	0.070	0.077	0.075
2	-0.157	-0.030	0.082	0.067	-0.093	0.076	0.074	0.066	0.072	0.066
3	-0.048	0.068	0.027	0.014	0.161	0.064	0.075	0.089	0.089	0.054
4	0.144	-0.031	-0.046	0.011	-0.063	0.082	0.087	0.080	0.081	0.074
5	0.125	0.113	-0.029	-0.113	-0.137	0.075	0.068	0.078	0.066	0.109
Mom	-0.745	-0.109	-0.094	0.138	0.340	0.158	0.090	0.068	0.062	0.090
Panel B: FF4										
Size/BM	1	2	3	4	5	1	2	3	4	5
1	-0.496	0.036	0.054	0.153	0.152	0.120	0.074	0.065	0.072	0.076
2	-0.103	0.002	0.095	0.068	-0.086	0.067	0.066	0.060	0.069	0.068
3	-0.005	0.096	0.040	0.025	0.195	0.065	0.076	0.078	0.086	0.862
4	0.143	-0.004	0.006	0.032	-0.015	0.076	0.075	0.078	0.077	0.075
5	0.147	0.096	-0.040	-0.087	-0.104	0.053	0.063	0.078	0.064	0.102
Mom	-0.016	0.259	-0.003	0.001	-0.046	0.091	0.057	0.074	0.067	0.062

Table 2. Parameters Estimates: Fama-French Model

The table reports posterior means and credibility intervals at the 95% level for a conditional version of the three factors Fama-French model. *Small-Growth* and *Small-Value* represent the first and fifth portfolios of the 25 Fama-French double-sorted portfolios. *Large-Growth* and *Large-Value* represent the 21st and the 25th double-sorted portfolios. *Loser* and *Winner* are the clustered portfolios in the lowest and highest quintiles sorted on past realized returns. The parameters are estimated through the Markov Chain Monte Carlo (MCMC) scheme described in the Appendix. The prior hyper-parameters are trained by using a ten-year pre-sample from 1963:07 to 1973:07. The testing sample is 1973:08-2013:01, monthly. Test portfolios and factors are described in the main text.

	Small-Growth	Small-Value	Large-Growth	Large-Value	Loser	Winner
Conditional Alphas						
δ_0^i	0.530	0.671	0.529	0.648	0.419	0.558
	[0.444 0.614]	[0.614 0.724]	[0.408 0.645]	[0.583 0.713]	[0.341 0.499]	[0.458 0.657]
$\bar{\alpha}_i$	-0.295	0.113	-0.304	0.105	-0.077	0.181
	[-0.494 -0.081]	[-0.029 0.256]	[-0.531 -0.075]	[-0.054 0.283]	[-0.243 0.074]	[-0.008 0.355]
τ_{i0}	1.544	0.745	1.447	1.017	1.511	0.927
	[1.395 1.742]	[0.679 0.827]	[1.235 1.747]	[0.921 1.111]	[1.378 1.654]	[0.834 1.051]
Market Betas						
δ_1^i	0.956	0.963	0.932	0.963	0.951	0.942
	[0.932 0.979]	[0.947 0.981]	[0.901 0.965]	[0.941 0.984]	[0.921 0.996]	[0.917 0.969]
$\bar{\beta}_i$	1.070	0.926	1.149	0.909	1.036	1.036
	[0.998 1.141]	[0.863 0.985]	[0.939 1.322]	[0.859 0.951]	[0.900 1.170]	[0.916 1.148]
τ_{i1}	0.020	0.027	0.149	0.016	0.068	0.081
	[0.019 0.043]	[0.017 0.036]	[0.108 0.219]	[0.011 0.023]	[0.056 0.117]	[0.053 0.103]
SMB Betas						
δ_2^i	0.960	0.946	0.920	0.971	0.935	0.946
	[0.935 0.981]	[0.921 0.968]	[0.879 0.957]	[0.955 0.986]	[0.894 0.969]	[0.917 0.971]
$\bar{\beta}_{i2}$	1.226	-0.248	0.201	1.010	-0.103	0.274
	[1.101 1.305]	[-0.321 -0.183]	[-0.021 0.441]	[0.906 1.107]	[-0.204 0.011]	[0.139 0.409]
τ_{i2}	0.041	0.039	0.205	0.029	0.058	0.077
	[0.023 0.068]	[0.016 0.065]	[0.131 0.271]	[0.021 0.045]	[0.033 0.088]	[0.061 0.129]
HML Betas						
δ_3^i	0.924	0.946	0.828	0.937	0.948	0.864
	[0.871 0.972]	[0.921 0.970]	[0.770 0.891]	[0.904 0.967]	[0.909 0.979]	[0.817 0.910]
$\bar{\beta}_{i3}$	-0.207	-0.499	0.033	0.571	0.797	-0.051
	[-0.311 -0.101]	[-0.587 -0.429]	[-0.168 0.256]	[0.459 0.671]	[0.672 0.924]	[-0.239 0.131]
τ_{i3}	0.061	0.047	0.461	0.073	0.051	0.302
	[0.039 0.118]	[0.024 0.069]	[0.369 0.617]	[0.025 0.115]	[0.041 0.078]	[0.254 0.356]
Idiosyncratic Risks						
δ_σ^i	0.958	0.957	0.947	0.939	0.940	0.949
	[0.852 0.992]	0.856 0.993]	[0.832 0.992]	[0.796 0.990]	[0.825 0.993]	[0.830 0.994]
σ_i^2	0.594	1.286	0.803	0.288	0.642	0.438
	[0.298 3.718]	[0.712 6.427]	[0.371 2.715]	[0.116 0.551]	[0.387 1.661]	[0.123 2.611]
τ_σ^i	0.145	0.143	0.158	0.139	0.156	0.143
	[0.114 0.233]	[0.112 0.225]	[0.122 0.253]	[0.113 0.224]	[0.119 0.246]	[0.117 0.226]

Table 3. Conditional Alphas

The table reports both the average and the in-sample (average) standard error of conditional alphas computed from different factor models on monthly excess returns of size, value, and momentum portfolios. Conditional estimates are made through a Markov Chain Monte Carlo (MCMC) scheme as described in the Appendix. The prior hyper-parameters are trained by using a ten-year pre-sample from 1963:07 to 1973:07, and the testing sample is 1973:08-2013:01. The test portfolios and factors are described in the main text. Top panel reports the results from the conditional CAPM. Middle panel reports the results computed from a standard Fama-French three factor model. Bottom panel reports the results from a four-factor model including momentum.

	Mean					St. Err				
CAPM										
Size/BM	1	2	3	4	5	1	2	3	4	5
1	-0.387	0.240	0.319	0.509	0.643	2.531	2.264	1.975	2.026	2.451
2	-0.083	0.173	0.386	0.431	0.441	1.997	1.655	1.544	1.657	2.055
3	-0.030	0.280	0.282	0.303	0.601	1.556	1.324	1.422	1.405	1.839
4	0.070	0.128	0.224	0.257	0.336	1.323	1.139	1.265	1.286	1.570
5	-0.095	0.082	0.022	0.054	0.146	1.074	0.889	1.157	1.059	1.705
Mom	-0.471	-0.010	-0.013	0.111	0.276	1.912	1.079	0.744	0.687	1.261
Fama-French 3 Factor										
Size/BM	1	2	3	4	5	1	2	3	4	5
1	-0.515	0.013	0.024	0.178	0.183	1.425	0.888	0.747	0.821	0.949
2	-0.123	-0.043	0.091	0.080	-0.061	0.876	0.761	0.698	0.820	0.838
3	0.007	0.105	0.029	-0.026	0.165	0.890	0.828	0.832	0.844	1.168
4	0.150	0.013	0.012	-0.028	0.008	0.823	0.900	0.930	0.951	1.120
5	0.171	0.099	-0.044	-0.115	-0.126	0.680	0.836	0.988	0.781	1.249
Mom	-0.408	-0.026	-0.020	0.038	0.209	1.165	0.684	0.684	0.433	0.789
Four Factors Model										
Size/BM	1	2	3	4	5	1	2	3	4	5
1	-0.532	-0.007	0.040	0.184	0.197	1.360	0.849	0.738	0.830	0.943
2	-0.112	-0.044	0.099	0.062	-0.070	0.811	0.746	0.697	0.809	0.841
3	-0.012	0.107	0.041	0.018	0.170	0.859	0.837	0.826	0.846	1.169
4	0.146	0.047	0.052	0.005	-0.004	0.814	0.885	0.879	0.919	1.114
5	0.190	0.101	-0.079	-0.110	-0.120	0.598	0.812	0.869	0.772	1.124
Mom	0.035	0.289	0.027	-0.109	-0.101	0.890	0.624	0.594	0.243	0.551

Table 4. Model-Implied Unconditional Alphas

The table reports the unconditional alphas and betas implied by the model dynamics. Conditional estimates are made through a Markov Chain Monte Carlo (MCMC) scheme as described in the Appendix. The prior hyper-parameters are trained by using a ten-year pre-sample from 1963:07 to 1973:07, and the testing sample is 1973:08-2013:01. The test portfolios and factors are described in the main text. Top panel reports the results from the conditional CAPM. Middle panel reports the results computed from a standard Fama-French three factor model. Bottom panel reports the results from a four-factor model including momentum. Bold-faced values denote estimates statistically significant at the 5% confidence level.

	Mean					St. Err				
CAPM										
Size/BM	1	2	3	4	5	1	2	3	4	5
1	-0.275	0.166	0.229	0.342	0.410	0.196	0.146	0.129	0.145	0.177
2	-0.060	0.116	0.244	0.256	0.260	0.127	0.113	0.108	0.122	0.151
3	-0.021	0.176	0.167	0.176	0.356	0.112	0.105	0.112	0.111	0.130
4	0.047	0.094	0.126	0.153	0.235	0.099	0.101	0.111	0.099	0.125
5	-0.055	0.052	0.020	0.034	0.083	0.094	0.091	0.093	0.088	0.113
Mom	-0.264	-0.004	-0.014	0.071	0.194	0.135	0.096	0.091	0.092	0.102
Fama-French Model										
Size/BM	1	2	3	4	5	1	2	3	4	5
1	-0.295	0.011	0.017	0.108	0.105	0.123	0.089	0.082	0.087	0.100
2	-0.092	-0.034	0.072	0.058	-0.044	0.083	0.073	0.078	0.087	0.082
3	0.003	0.073	0.025	-0.014	0.069	0.081	0.086	0.072	0.086	0.100
4	0.101	0.007	0.008	-0.022	0.006	0.083	0.081	0.088	0.084	0.087
5	0.113	0.054	-0.029	-0.076	-0.077	0.085	0.082	0.089	0.076	0.095
Mom	-0.304	-0.023	-0.019	0.039	0.181	0.141	0.086	0.083	0.077	0.108
Four Factors										
Size/BM	1	2	3	4	5	1	2	3	4	5
1	-0.348	-0.007	0.026	0.104	0.116	0.130	0.084	0.076	0.087	0.100
2	-0.088	-0.030	0.075	0.048	-0.051	0.076	0.081	0.073	0.088	0.081
3	-0.010	0.075	0.030	0.014	0.086	0.079	0.090	0.076	0.084	0.097
4	0.106	0.036	0.032	0.003	-0.005	0.085	0.082	0.083	0.077	0.083
5	0.139	0.052	-0.050	-0.084	-0.080	0.084	0.087	0.088	0.075	0.101
Mom	0.026	0.199	0.013	-0.092	-0.077	0.094	0.091	0.085	0.072	0.078

Table 5. Market Betas and Standard Predictors

The table reports the slope estimates when the predictive betas are regressed on a set of standard state variables. The predictive market betas are computed integrating out both investors' uncertainty on the structural parameters of the three-factor Fama-French model. The sample is 1973:01-2013:12. The state variables are the dividend yield (dy), the earnings-to-price ratio (ep), the dividend-payout ratio (dpr), net equity expansion (ntis), the default spread (def), log inflation (inf), industrial production (ip), the year-on-year consumption growth (cons), the M2 monetary aggregate (m2), the lagged excess return on the market ($ret_{(-1)}$), and the term spread (term). Bold-faced values denote estimates statistically significant at the 5% confidence level. Standard errors are corrected for autocorrelation and heteroschedasticity. The test portfolios and factors are described in the main text.

	Small-Growth	Small-Value	Large-Growth	Large-Value	Loser	Winner
Predictor						
dy	0.885	0.205	-0.117	-0.191	-0.044	0.316
ep	0.890	0.750	-0.127	-0.995	0.302	-0.168
dpr	-0.391	0.553	0.128	-0.441	-0.066	-0.205
def	0.016	-0.125	-0.206	-0.134	-0.073	0.049
ntis	-0.023	0.237	-0.096	-0.274	-0.261	0.102
inf	0.061	-0.198	0.092	0.096	0.136	-0.030
ip	0.008	-0.119	-0.224	-0.004	0.093	0.162
cons	-0.120	0.637	0.545	0.220	-0.859	0.231
m2	-0.115	0.142	-0.222	-0.202	-0.152	-0.342
ret(-1)	-0.085	0.026	-0.075	-0.043	-0.018	0.060
term	-0.194	0.040	0.038	0.183	0.270	-0.216
Adj R^2	0.483	0.261	0.386	0.231	0.513	0.321

Table 6. Variance Ratios

The table reports the results of variance decomposition test. $\mathcal{VR}1$ is the ratio of the variance of a model predicted returns and the variance of expected returns estimated from a projection on a set of instruments \mathbf{Z}_t . $\mathcal{VR}2$ is the ratio of the variance of the predictable part of returns not explained by a model and the variance of projected returns. The instrumental variables are the lagged monthly dividend yield on the NYSE/AMEX, the lagged yield of a Baa corporate bond, and the lagged spread of long- vs. short-term government bond yields. Boldfaced numbers indicate the highest $\mathcal{VR}1$ and the lowest $\mathcal{VR}2$ respectively.

	CAPM						Fama-French						Four-Factor Model					
	$\mathcal{VR}1$			$\mathcal{VR}2$			$\mathcal{VR}1$			$\mathcal{VR}2$			$\mathcal{VR}1$			$\mathcal{VR}2$		
	2.5%	50%	97.5%	2.5%	50%	97.5%	2.5%	50%	97.5%	2.5%	50%	97.5%	2.5%	50%	97.5%	2.5%	50%	97.5%
S1B1	0.480	0.790	1.113	0.219	0.400	0.602	0.511	0.775	1.073	0.039	0.453	0.453	0.461	0.701	0.915	0.247	0.467	0.676
S1B2	0.311	0.545	0.789	0.382	0.552	0.762	0.431	0.645	0.904	0.089	0.467	0.467	0.403	0.587	0.757	0.219	0.404	0.599
S1B3	0.234	0.444	0.672	0.479	0.650	0.853	0.371	0.561	0.775	0.133	0.505	0.505	0.363	0.535	0.693	0.217	0.415	0.604
S1B4	0.156	0.367	0.569	0.398	0.539	0.721	0.329	0.495	0.682	0.061	0.360	0.360	0.340	0.508	0.658	0.065	0.224	0.377
S1B5	0.039	0.182	0.340	0.271	0.368	0.508	0.225	0.345	0.482	-0.017	0.194	0.194	0.227	0.330	0.423	-0.094	0.015	0.118
S2B1	0.529	0.846	1.167	0.265	0.434	0.655	0.554	0.833	1.166	0.002	0.398	0.398	0.525	0.793	1.027	0.184	0.404	0.611
S2B2	0.435	0.702	0.990	0.307	0.458	0.658	0.553	0.835	1.163	-0.027	0.358	0.358	0.548	0.805	1.047	0.098	0.314	0.518
S2B3	0.165	0.340	0.528	0.362	0.510	0.666	0.329	0.512	0.720	0.045	0.358	0.358	0.341	0.518	0.687	0.079	0.248	0.419
S2B4	0.192	0.373	0.575	0.353	0.502	0.678	0.395	0.597	0.831	0.032	0.333	0.333	0.403	0.585	0.767	0.009	0.176	0.339
S2B5	0.343	0.568	0.818	0.281	0.422	0.591	0.462	0.696	0.963	-0.020	0.271	0.271	0.456	0.671	0.860	-0.050	0.097	0.250
S4B1	0.257	0.449	0.643	0.206	0.329	0.480	0.468	0.701	0.957	-0.007	0.292	0.292	0.453	0.679	0.858	0.014	0.182	0.365
S4B2	0.328	0.563	0.812	0.280	0.434	0.624	0.590	0.898	1.269	0.019	0.391	0.391	0.568	0.860	1.134	0.092	0.326	0.575
S4B3	0.447	0.702	0.967	0.387	0.567	0.777	0.740	0.906	1.525	0.045	0.505	0.505	0.713	1.030	1.340	0.172	0.436	0.730
S4B4	0.269	0.469	0.663	0.285	0.423	0.577	0.592	0.896	1.231	0.087	0.487	0.487	0.601	0.884	1.142	0.083	0.320	0.557
S4B5	0.386	0.637	0.912	0.432	0.624	0.841	0.587	0.918	1.252	0.191	0.639	0.639	0.603	0.881	1.148	0.174	0.419	0.677
M1	0.722	1.119	1.547	0.480	0.710	0.992	0.709	0.852	1.418	0.166	0.281	0.481	0.578	0.839	1.097	0.604	0.900	1.180
M2	0.340	0.584	0.819	0.371	0.546	0.744	0.522	0.800	1.106	0.097	0.496	0.496	0.479	0.704	0.906	0.218	0.423	0.639
M3	0.338	0.556	0.773	0.400	0.542	0.721	0.505	0.789	1.119	0.056	0.462	0.462	0.522	0.782	1.021	0.085	0.292	0.508
M4	0.296	0.516	0.748	0.570	0.773	0.996	0.545	0.852	1.246	0.150	0.657	0.657	0.608	0.926	1.236	0.181	0.452	0.728
M5	0.252	0.479	0.734	0.320	0.500	0.700	0.442	0.759	1.086	-0.032	0.337	0.337	0.632	0.934	1.175	-0.120	0.078	0.259

Table 7. Marginal Likelihoods and Bayes Factors

The table reports the values of the (log)marginal likelihoods and the relative (log of) Bayes Factors for different factor model specifications. $\mathcal{BF}_{i,j}$ is the Bayes Factor of the model \mathcal{M}_i against \mathcal{M}_j . Here \mathcal{M}_0 is the CAPM, \mathcal{M}_1 the three-factor Fama-French model, and \mathcal{M}_2 the four-factor model proposed in Charhart (1997). The values reported are also disaggregated by computing the contributions from each of the portfolios under investigation. The prior hyper-parameters are trained by using a ten-year pre-sample from 1963:07 to 1973:07. The testing sample is 1973:08-2013:01, monthly. The test portfolios and factors are described in the main text. Bold-faced values denote the model with the highest log-marginal likelihood.

	\mathcal{M}_0	\mathcal{M}_1	\mathcal{M}_2	$\mathcal{BF}_{2,0}$	$\mathcal{BF}_{1,2}$	$\mathcal{BF}_{1,0}$
S1B1	-1382.30	-1001.00	-1066.49	631.62	130.97	762.59
S1B2	-1339.79	-852.35	-867.33	944.92	29.95	974.87
S1B3	-1267.31	-785.27	-798.78	937.06	27.01	964.06
S1B4	-1250.19	-812.80	-819.48	861.42	13.37	874.79
S1B5	-1297.61	-821.77	-849.59	896.04	55.65	951.69
S2B1	-1337.71	-831.16	-856.54	962.34	50.77	1013.11
S2B2	-1227.59	-785.30	-801.33	852.52	32.05	884.57
S2B3	-1169.87	-724.53	-763.21	813.31	77.37	890.68
S2B4	-1159.27	-781.80	-801.63	715.28	39.65	754.93
S2B5	-1226.42	-869.71	-870.37	712.10	1.33	713.43
S4B1	-1092.27	-833.91	-846.94	490.68	26.05	516.72
S4B2	-959.37	-871.26	-879.87	159.00	17.20	176.20
S4B3	-977.14	-829.26	-880.34	193.60	102.17	295.76
S4B4	-1067.15	-941.72	-957.14	220.01	30.84	250.85
S4B5	-1140.97	-1054.05	-1052.44	177.07	-3.22	173.85
S5B1	-962.18	-591.17	-669.27	585.82	156.19	742.01
S5B2	-883.01	-848.57	-865.00	36.01	32.86	68.87
S5B3	-1014.16	-851.39	-919.48	189.36	136.18	325.54
S5B4	-1036.92	-847.66	-859.51	354.82	23.71	378.52
S5B5	-1258.19	-1007.41	-1053.77	408.85	92.72	501.56
M1	-1236.60	-863.25	-1016.82	439.56	307.13	746.68
M2	-937.22	-652.03	-810.77	252.92	317.46	570.38
M3	-750.44	-605.64	-674.14	152.60	137.02	289.61
M4	-757.52	-443.96	-553.23	408.59	218.53	627.12
M5	-1066.18	-581.70	-786.76	558.84	410.14	968.97
Global	-33478.44	-24564.93	-25851.30	15254.27	2572.74	17827.01

Figure 1. Dynamic Hypothesis Testing: Joint Conditional Alphas

This figure plots the posterior probability of the null hypothesis that pricing errors are jointly not statistically different from zero at time t , given the whole sample information. Estimates are based on the Markov Chain Monte Carlo (MCMC) scheme (see the appendix). Prior hyper-parameters are trained by using a ten-year pre-sample from 1963:07 to 1973:07. The testing sample is 1973:08-2013:01, monthly. Factors and test portfolios are described in the main text.

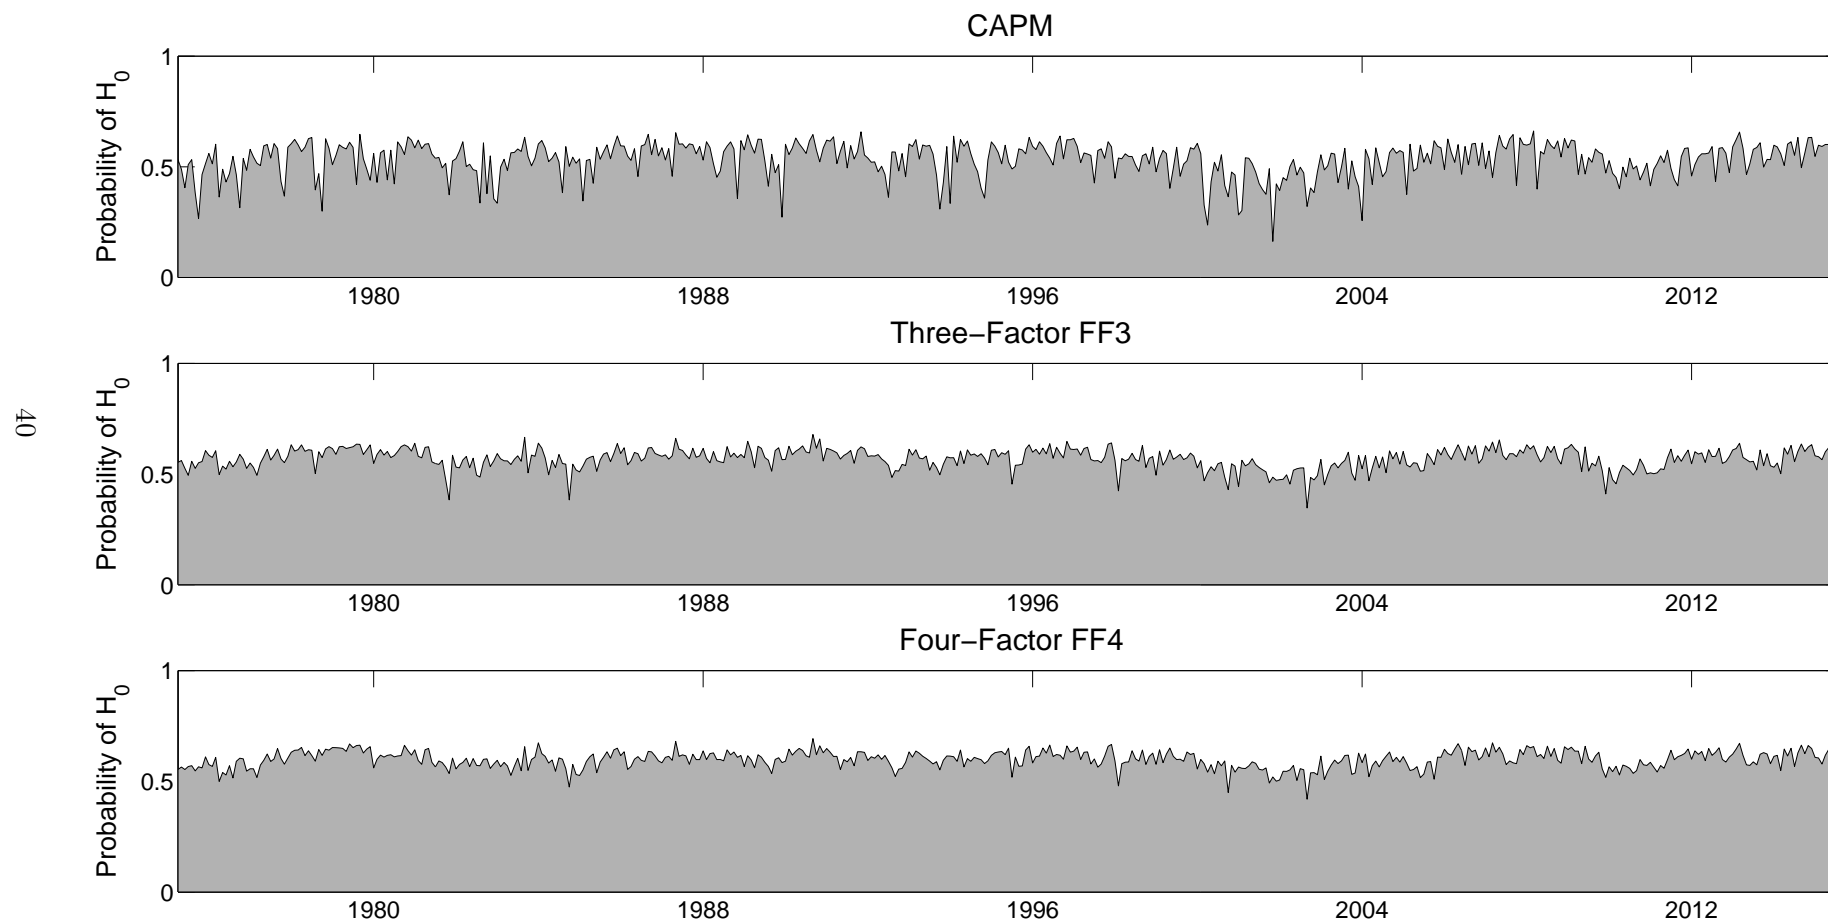


Figure 2. Dynamic Hypothesis Testing: Asset Specific Conditional Alphas (CAPM)

This figure plots the posterior probabilities of the null hypothesis that the conditional CAPM holds at time t , given the whole sample information. Estimates are based on the Markov Chain Monte Carlo (MCMC) scheme (see the appendix). Prior hyper-parameters are trained by using a ten-year pre-sample from 1963:07 to 1973:07. The testing sample is 1973:08-2013:01, monthly. The test portfolios are described in the main text.

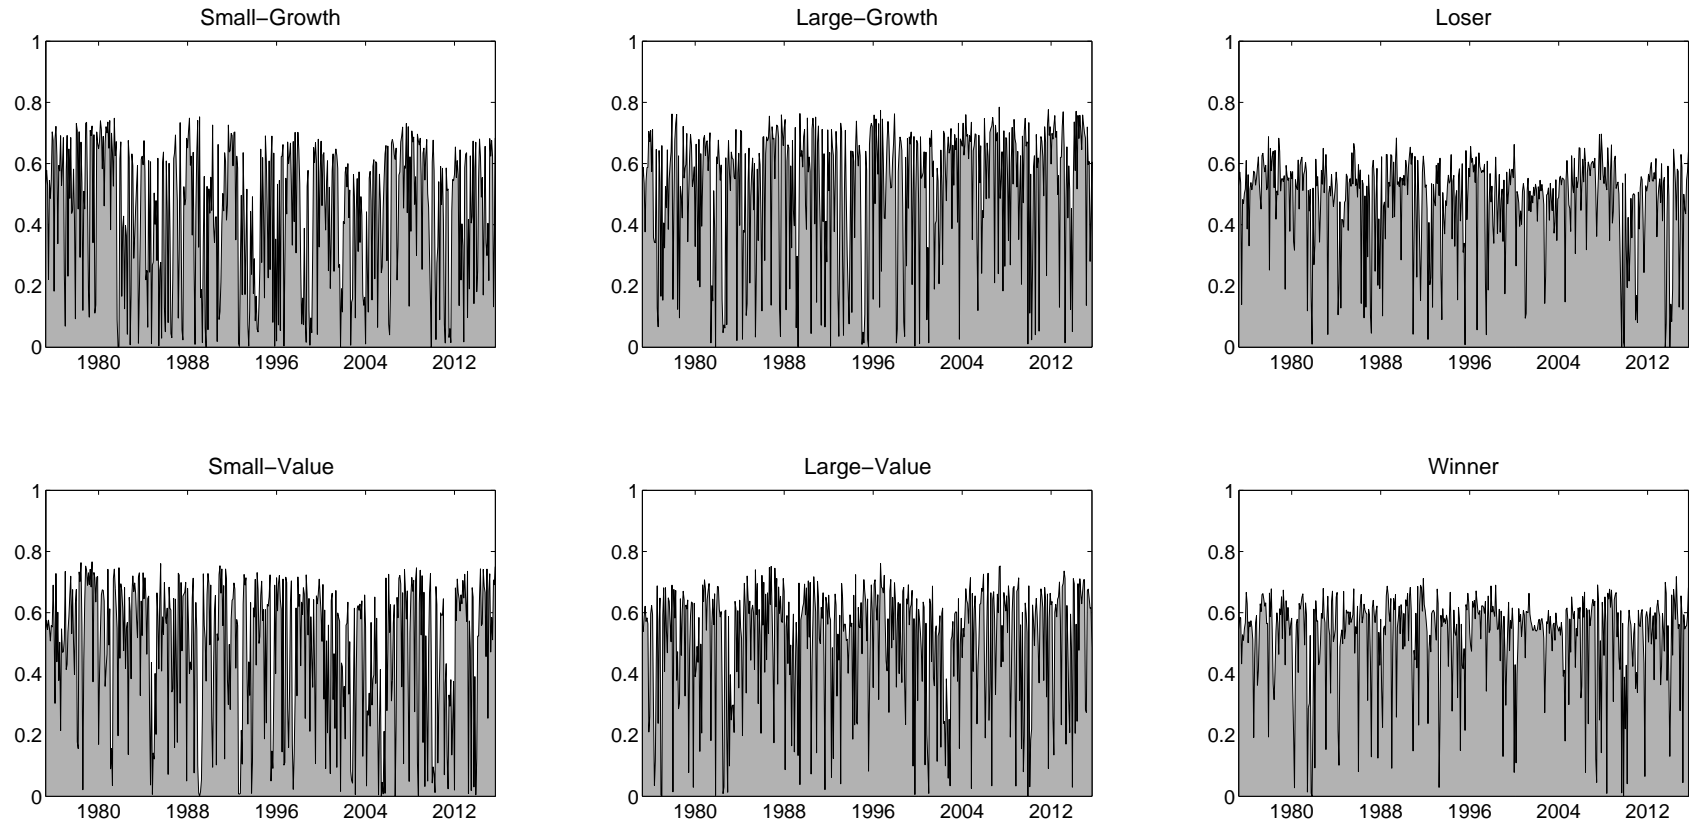


Figure 3. Dynamic Hypothesis Testing: Asset Specific Conditional Alphas (Three-Factor Model)

This figure plots the posterior probabilities of the null hypothesis that the three-factor Fama-French model holds at time t , given the whole sample information. Estimates are based on the Markov Chain Monte Carlo (MCMC) scheme (see the appendix). Prior hyper-parameters are trained by using a ten-year pre-sample from 1963:07 to 1973:07. The testing sample is 1973:08-2013:01, monthly. The test portfolios are described in the main text.

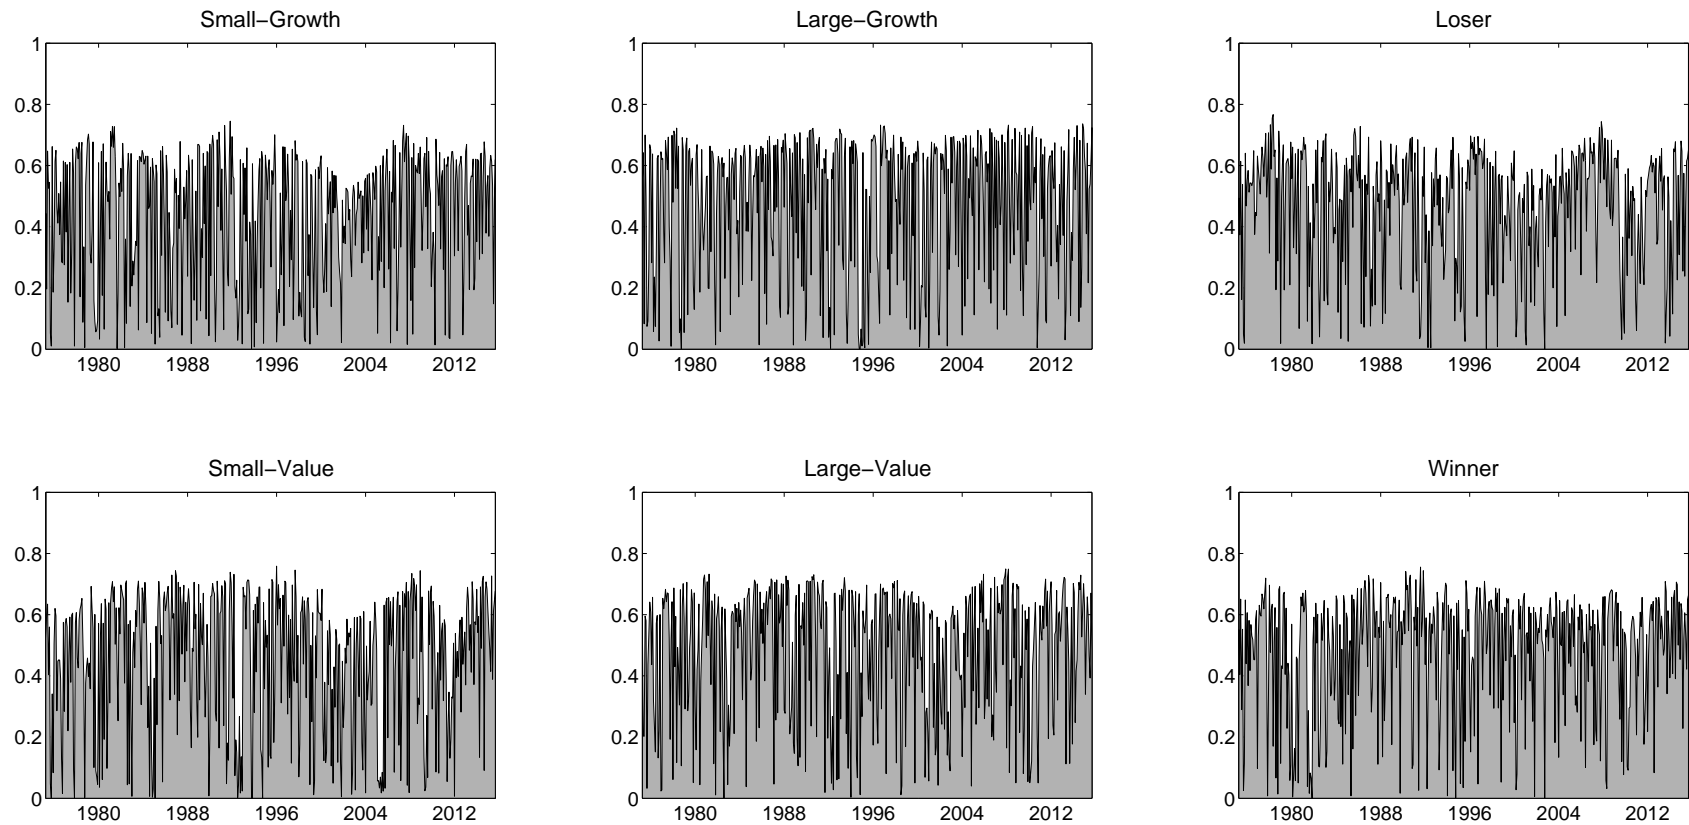


Figure 4. Dynamic Hypothesis Testing: Asset Specific Conditional Alphas (Four-Factor Model)

This figure plots the posterior probabilities of the null hypothesis that the four-factor model of Charhart (1997) holds at time t , given the whole sample information. Estimates are based on the Markov Chain Monte Carlo (MCMC) scheme (see the appendix). Prior hyper-parameters are trained by using a ten-year pre-sample from 1963:07 to 1973:07. The testing sample is 1973:08-2013:01, monthly. The test portfolios are described in the main text.

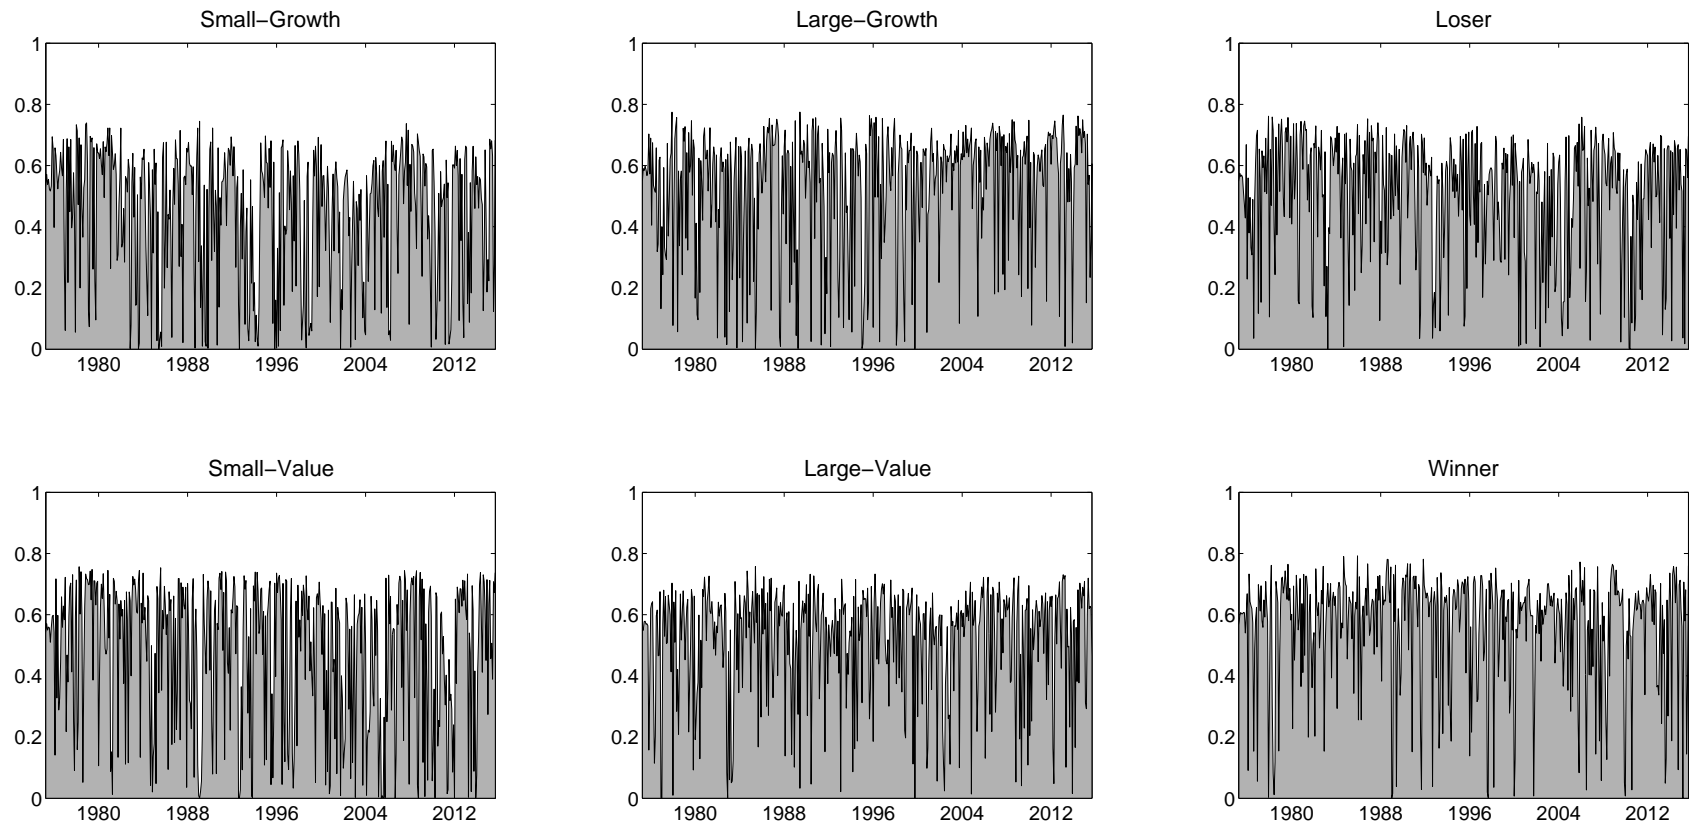


Figure 5. Market Risk Premium

This figure plots sequence of posterior distributions of the market risk premium and its conditional variance. Estimates are made through a Markov Chain Monte Carlo (MCMC) scheme. Prior hyper-parameters are trained by using a ten-year pre-sample from 1963:07 to 1973:07. The testing sample is 1973:08-2013:01, monthly. The market risk premium is computed from no-arbitrage restrictions (4). The dark line represent the median value while the shaded area is the 95% interval. The red line is the conditional volatility estimate from a GARCH(1,1) fitted on historical returns on the CRSP value-weighted index.

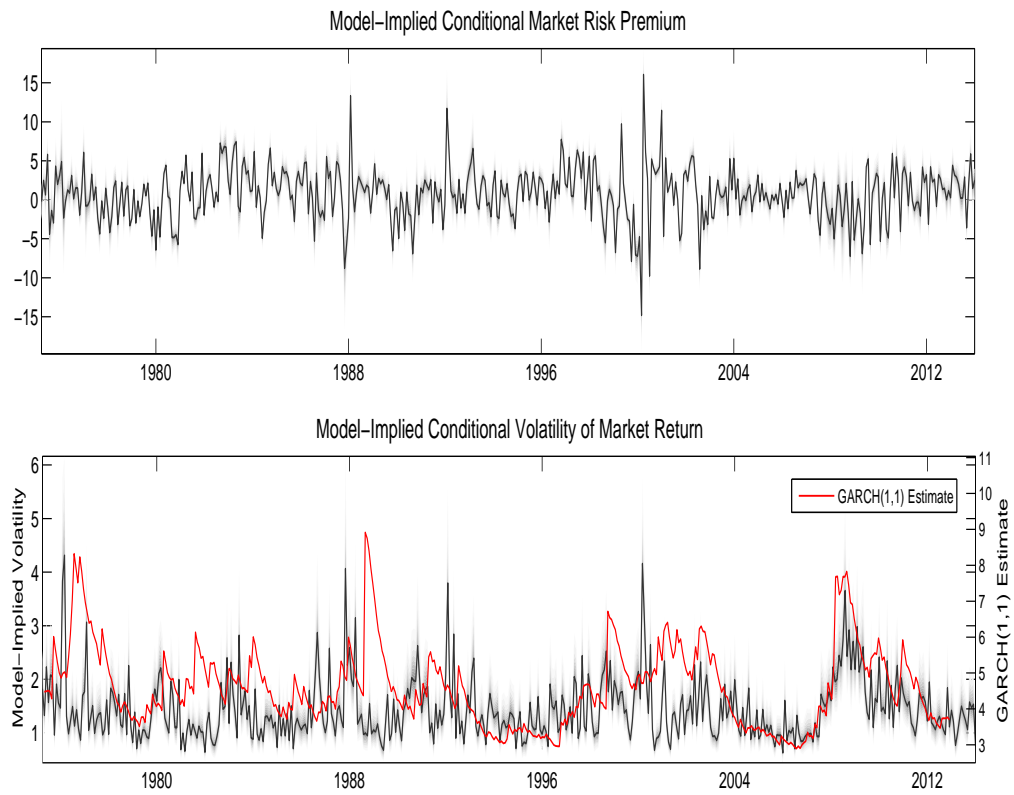


Figure 6. Dynamic Hypothesis Testing: Cross-Sectional Pricing Error

This figure plots the posterior probabilities of the null hypothesis that the cross-sectional pricing error is not statistically different from zero at time t . Estimates are made through a Markov Chain Monte Carlo (MCMC) scheme (see the appendix). Prior hyper-parameters are trained by using a ten-year pre-sample from 1963:07 to 1973:07. The testing sample is 1973:08-2013:01, monthly. The test portfolios and factors are described in the main text.

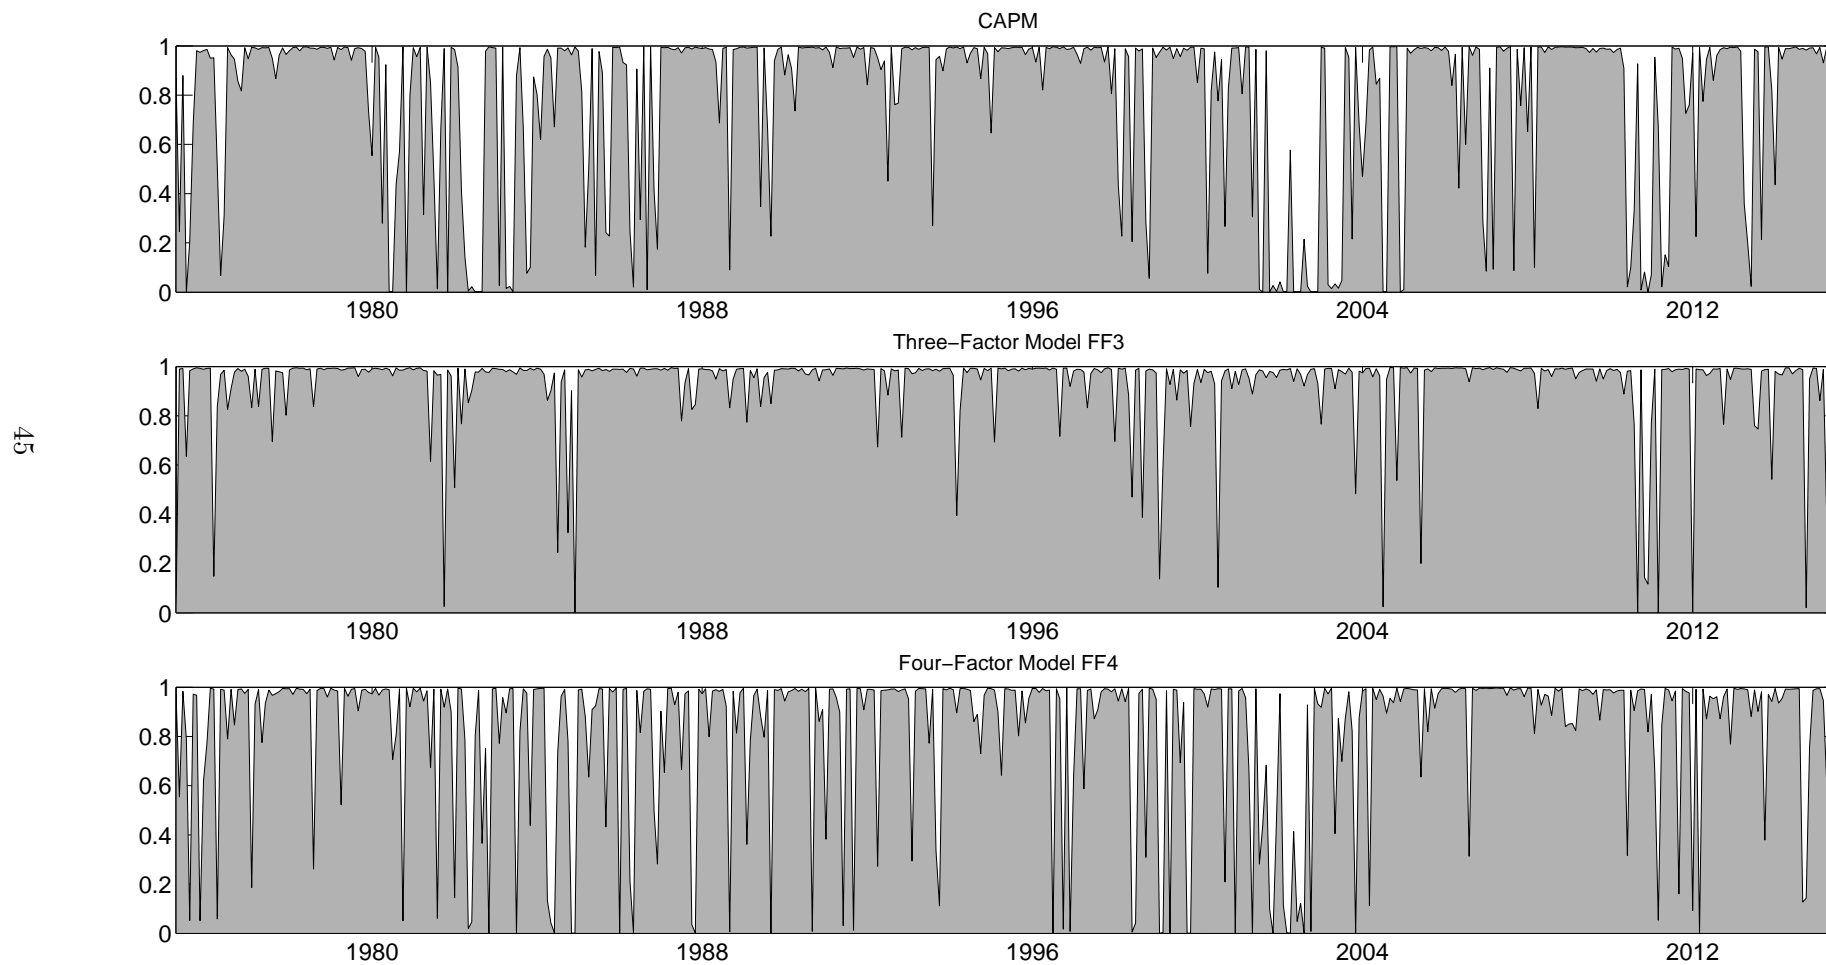


Figure 7. Conditional Market Betas

This figure plots sequence of posterior distributions of the market betas from the three-factor Fama-French model. Estimates are made through the Markov Chain Monte Carlo (MCMC) scheme (i.e. the model with full uncertainty). The prior hyper-parameters are trained by using a ten-year pre-sample from 1963:07 to 1973:07. The testing sample is 1973:08-2013:01, monthly. The test portfolios are described in the main text. The dark line represent the median value while the shaded area is the 95% interval.

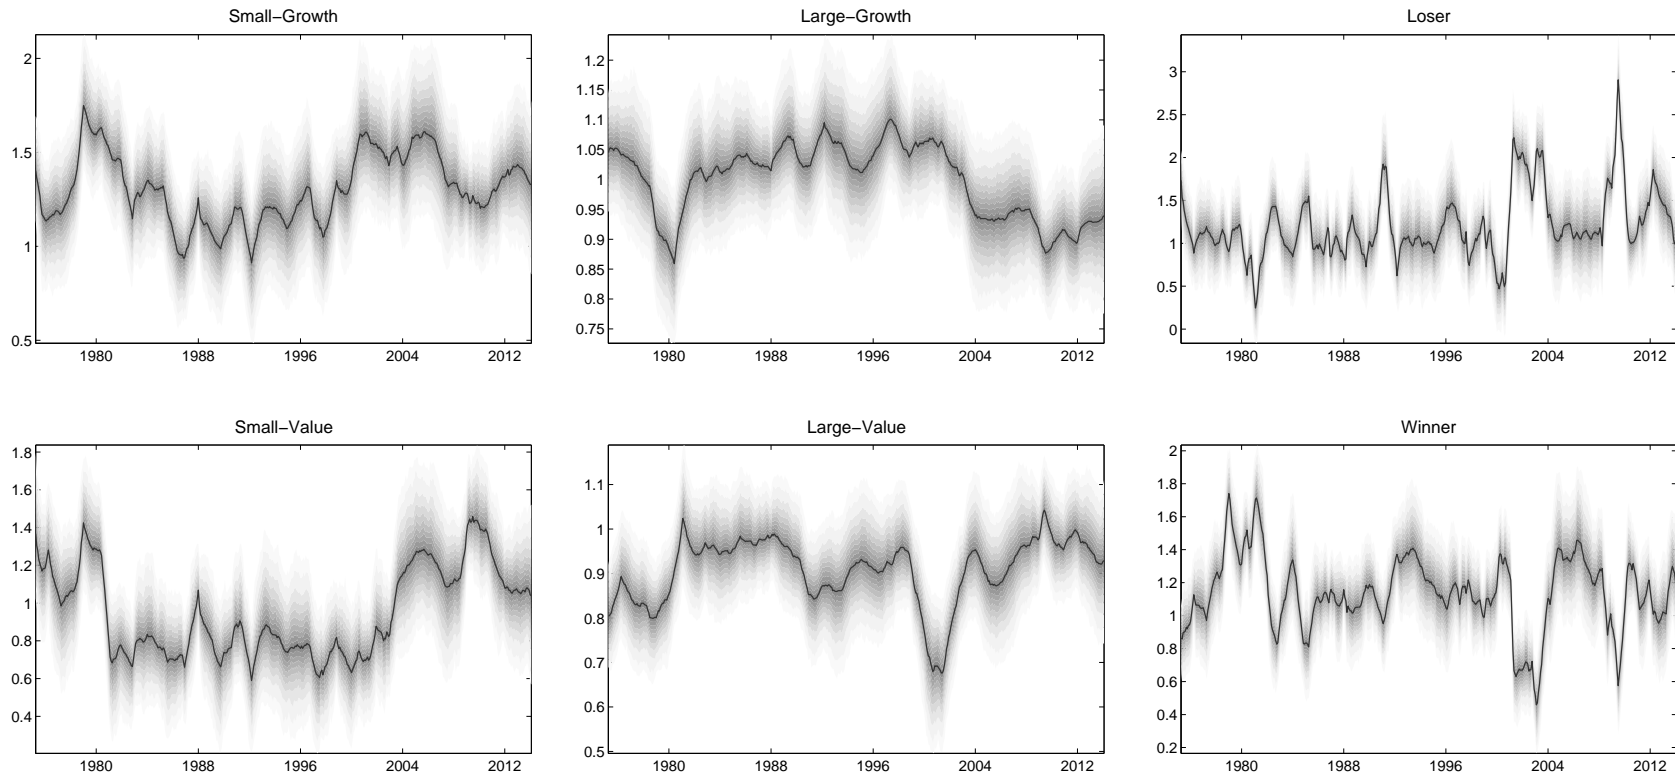


Figure 8. Conditional SMB Betas

This figure plots sequence of posterior distributions of the betas on the SMB factor from the three-factor Fama-French model. Estimates are made through the Markov Chain Monte Carlo (MCMC) scheme (i.e. the model with full uncertainty). The prior hyper-parameters are trained by using a ten-year pre-sample from 1963:07 to 1973:07. The testing sample is 1973:08-2013:01, monthly. The test portfolios are described in the main text. The dark line represent the median value while the shaded area is the 95% interval.

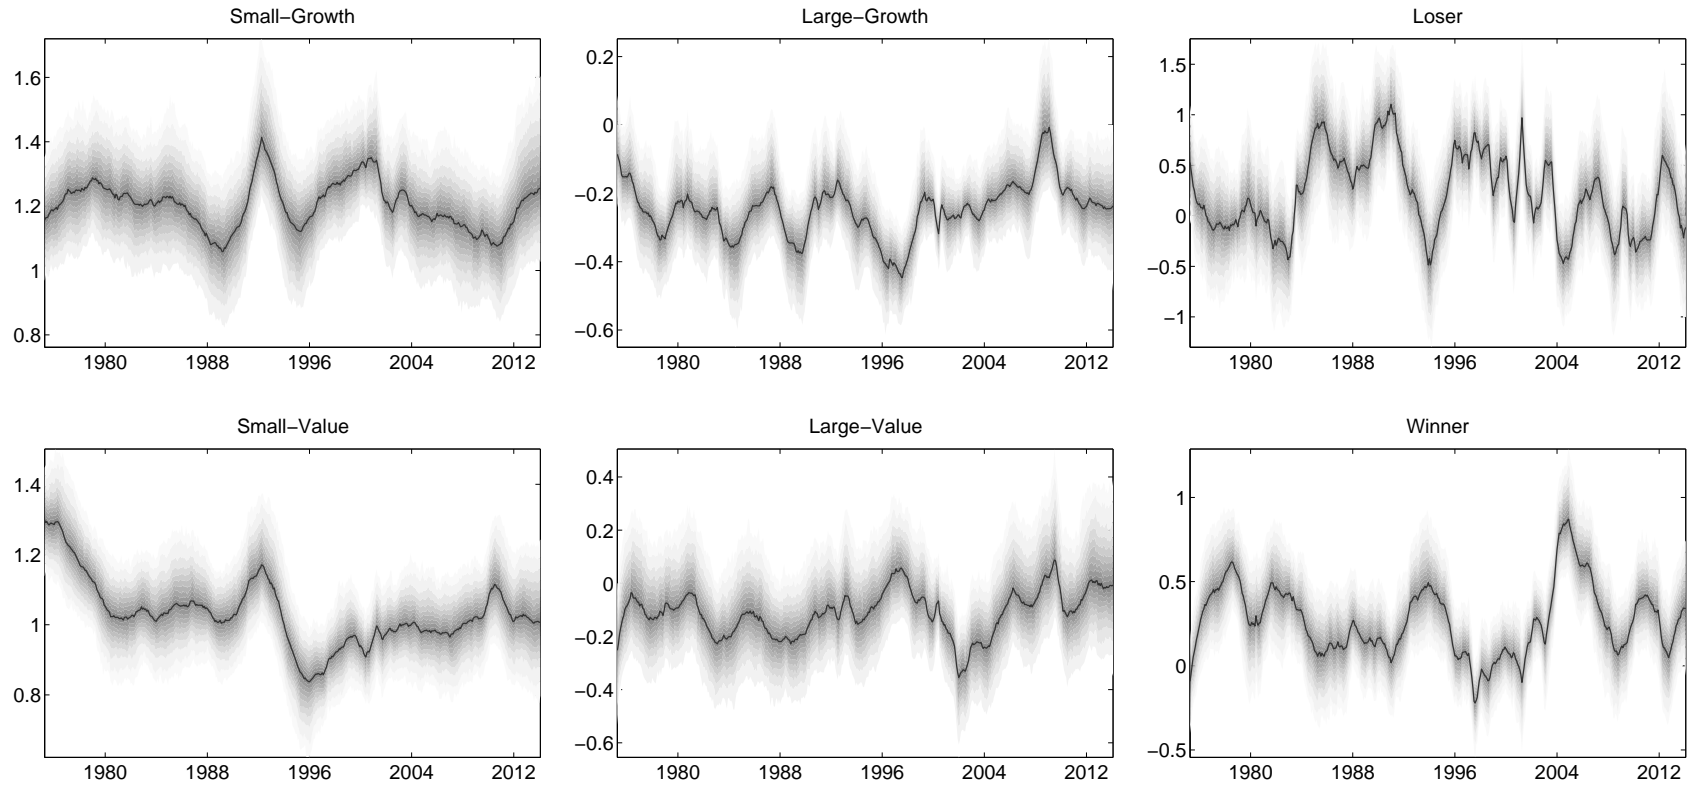


Figure 9. Idiosyncratic risks

This figure plots the sequence of posterior distributions of (the square root of) idiosyncratic risks $\sigma_{i,t}$. Estimates are made through the Markov Chain Monte Carlo (MCMC) scheme (i.e. the model with full uncertainty). The prior hyper-parameters are trained by using a ten-year pre-sample from 1963:07 to 1973:07. The testing sample is 1973:08-2013:01, monthly. The test portfolios are described in the main text. The dark line represent the median value while the shaded area is the 95% interval.

