2013 | 19

Working Paper

Norges Bank Research

Macroeconomic factors strike back: A Bayesian change-point model of time-varying risk exposures and premia in the U.S. cross-section

Daniele Bianchi, Massimo Guidolin and Francesco Ravazzolo

NB NORGES BANK

Working papers fra Norges Bank, fra 1992/1 til 2009/2 kan bestilles over e-post:
servicesenter@norges-bank.no

Fra 1999 og senere er publikasjonene tilgjengelige på www.norges-bank.no

Working papers inneholder forskningsarbeider og utredninger som vanligvis ikke har fått sin endelige form. Hensikten er blant annet at forfatteren kan motta kommentarer fra kolleger og andre interesserte. Synspunkter og konklusjoner i arbeidene står for forfatternes regning.

Working papers from Norges Bank, from 1992/1 to 2009/2 can be ordered by e-mail:
servicesenter@norges-bank.no

Working papers from 1999 onwards are available on www.norges-bank.no

Norges Bank's working papers present research projects and reports (not usually in their final form) and are intended inter alia to enable the author to benefit from the comments of colleagues and other interested parties. Views and conclusions expressed in working papers are the responsibility of the authors alone.

ISSN 1502-8143 (online)
ISBN 978-82-7553-772-8 (online)

Macroeconomic Factors Strike Back: A Bayesian Change-Point Model of Time-Varying Risk Exposures and Premia in the U.S. Cross-Section*

Daniele Bianchi †, Massimo Guidolin‡ and Francesco Ravazzolo§

October 17, 2013

Abstract

This paper proposes a Bayesian estimation framework for a typical multi-factor model with time-varying risk exposures to macroeconomic risk factors and corresponding premia to price U.S. publicly traded assets. The model assumes that risk exposures and idiosynchratic volatility follow a break-point latent process, allowing for changes at any point on time but not restricting them to change at all points. The empirical application to 40 years of U.S. data and 23 portfolios shows that the approach yields sensible results compared to previous two-step methods based on naive recursive estimation schemes, as well as a set of alternative model restrictions. A variance decomposition test shows that although most of the predictable variation comes from the market risk premium, a number of additional macroeconomic risks, including real output and inflation shocks, are significantly priced in the cross-section. A Bayes factor analysis massively favors of the proposed change-point model.

Key words and phrases: Structural breaks, Change-point model, Stochastic volatility, Multi-factor linear models, Asset Pricing.

JEL codes: G11, C53.

.

.

*We are grateful to Carlos Carvalho and to seminar participants at the ICEE 2013 in Genoa, Norges Bank, the SIS 2013 in Brescia, the First Vienna workshop on high-dimensional time series in macroeconomics and finance, IHS 2013, the 6th International Conference on Computational and Financial Econometrics in Oviedo (CFE 12), and University of St. Gallen, for helpful comments. The views expressed in this paper are our own and do not necessarily reflect those of Norges Bank.

†Bocconi University. E-mail: daniele.bianchi@unibocconi.it.

‡IGIER, Bocconi University. E-mail: massimo.guidolin@unibocconi.it.

§Norges Bank and BI Norwegian Business School. E-mail: Francesco.Ravazzolo@Norges-Bank.no.

1. Introduction

Can a selected set of macroeconomic variables explain the cross-sectional behavior of U.S. stock and bond returns, i.e., why different assets earn different average rates of return? This simple question lies at the heart of the burgeoning field of macro-finance. Remarkably enough, the answer provided by at least 20 years of research on this crucial question has been predominantly negative (see e.g., Chan, Karceski, and Lakonishok, 1998; McQueen and Roley, 1993; Shanken and Weinstein, 2006): although occasional nuances to this fundamentally negative result have been reported (e.g., Flannery and Protopapadakis, 2002; Kramer, 1994; McQueen and Roley, 1993, conditioning on the state of the economy), it is common wisdom that macroeconomic factors can hardly explain the cross-sectional dynamics of asset valuations and returns of U.S. stock and bond portfolios. Such a disconnect between changes in aggregate variables representing sources of systematic risk—like in the case of output and inflation growth news—and asset returns has long represented a puzzle.

In this paper we propose and estimate through Bayesian methods a flexible parametric multi-factor, stochastic volatility asset pricing model in which both risk exposures (betas) and the prices of a number of macroeconomic risk factors are time-varying and effectively explain the cross-section of U.S. stock and bond returns (see Gungor and Luger, 2013). Time variation is modelled as a latent, change-point process. We show that an explicit parameterization of latent change-points in betas and risk premia plays a dominant role. By comparing our baseline model with restricted versions of the same, we also provide evidence that both stochastic volatility and infrequent but possibly large parameter instability are key drivers of the capability of the model to capture cross-sectional return dynamics.

Drawing a precisely estimated link between time-varying betas on selected macroeconomic risk factors and stock and bond excess returns also speaks to the very heart of finance theory, because any evidence uncovered bears on the fundamental issue of the key features of the general pricing mechanism, called the stochastic discount factor (SDF), underlying observed security prices. Practically, the SDF identifies a change of measure from the objective (physical) probability to a risk-neutralized one that allows us to price all assets and portfolios by simply discounting to the present the expected stream of future of payoffs that they will produce before their expiration date. Such a change of measure depends on the shape of the (aggregate) risk aversion function of investors and therefore reflects the way in which systematic risk factors are priced in the aggregate (see e.g., Cochrane, 2005; Singleton, 2006). In our paper we show that is both possible and useful to connect such an SDF (assumed to exist and to be unique) and macroeconomic risks.

A related question concerns the most appropriate methods available to researchers to learn about such the SDF that underlies the cross-section of asset returns. Our paper offers a contribution to an extensive literature on the estimation of empirical SDFs, specializing to a particular set of linear multi-factor models, offers a novel statistical framework to implement such models, and shows how this works using an empirically relevant application. With reference to an application to 40 years of monthly data on excess returns on 23 key portfolios of securities traded in the U.S., we show that while commonly used methods to estimate macro-based linear factor models fail to lead to sensible conclusions, an encompassing Gibbs sampling algorithm that allows for instability in factor exposures and risk premia using break-point processes delivers encouraging results.

Following the seminal work of Fama and MacBeth (1973), two-step multi-factor asset pricing models (MFAPMs) have been commonly used to estimate multi-factor models. Fama-MacBeth's (henceforth F-MB) approach, first proposed for the plain vanilla CAPM but then extended to a wider class of linear models, corresponds to a very simple algorithm: the risk premium on any asset or portfolio is decomposed as the sum of risk exposures to a number of risk factors multiplied by the associated unit price for each factor. The algorithm uses a first set of rolling window, time series regressions to obtain estimates of the betas, followed by a second-pass set of cross-sectional (across assets) regressions that using the first-pass risk exposures as inputs to derive time-varying estimates of the premia. The limitations of this methodology are now well-understood:[1] most inferential statements made as a result of the second-pass would be valid if and only if one could assume that the first-pass betas were fixed in repeated samples, which contradicts their random sample nature deriving from their being least squares estimates. Unless additional assumptions are introduced, this creates a problem with generated regressors being used in the second-step, which makes most of the inferential statements commonly made when the resulting error-in-variables problems are ignored invalid (see Pagan, 1984).[2] F-MB's approach also suffers from another problem: although identifying time-variation in risk exposures and premia with a rolling window least square estimation is robust because it is nonparametric, the length of the window is usually chosen in an arbitrary way and this can result in a severe loss of efficiency (see e.g., Maheu and McCurdy, 2009).

To overcome these problems, we introduce a different approach where time variation in risk exposures and premia is explicitly modelled as a break-point process. Specifically, we model risk exposures as latent stochastic processes in a mixture innovation framework as in Giordani, Kohn and van Dijk (2007), Giordani and Kohn (2008), Groen, Paap, and Ravazzolo (2013), Maheu and Gordon (2008). The parameters of interest are constant unless a break-point variable takes a unit value, in which case the parameters are allowed to jump to a new level, as a result of a normally distributed shock (see Jostova and Philipov, 2005). Furthermore, to consistently overcome the problems with generated regressors, the model is estimated in a single step by using a Bayesian approach, following the seminal work by McCulloch and Rossi's (1991) and Geweke and Zhou (1996).[3] In this paper, we provide an exact finite

[1]Our paper is not about how to produce better standard errors than under F-MB's methods in asset pricing tests involving data panels, to take into account cross- and own-serial correlation effects. Petersen (2009) reviews these methods and performs thorough comparisons. Geweke and Zhou (1996) discuss the difference between the two endevours.

[2]In practice, the classical two-step procedure either does not provide a known asymptotic distribution for functions of interest in applied asset pricing work or these asymptotic distributions may not be reliable in finite samples, even when they are available (see Petersen, 2009).

[3]However, McCulloch and Rossi's approach remains a two-pass procedure in which the factors are extracted before

sample statistical framework for testing multi-factor models. By construction, our approach represents a single-step procedure that yields exact inferences; given the fact that there are unobservable factors in the assumed return generating process, our framework implicitly incorporates this uncertainty into our inferences. Moreover, our approach makes it possible to compute the posterior distribution of virtually any function of the parameters that can be useful to implement economic tests (e.g., variance ratio and decomposition tests applied to the predictable portions of asset returns).[4]

Our main results can be summarized as follows. First, using a variety of metrics—such as Bayesian factors and average pricing error performance—we obtain evidence of the importance of capturing both instability in betas and in stochastic volatility; additionally, simpler time-varying parameter models in which betas follow random walk processes in which breaks are frequent but of modest size appear to be outperformed by our change-point model. The Bayesian (posterior median) estimates of the risk premia are stable over time and a few of them are precisely estimated. Moreover, a variance decomposition test shows that the by considering model instability, together with parameter uncertainty, the amount of cross-sectional excess return variation explained by the factor model increases with respect not only to a standard F-MB, but also with respect to the case in which specific parsimonious restrictions on the dynamics of both factor sensitivities and idiosyncratic risks are imposed. Second, the Bayesian time-varying betas, stochastic volatility model leads to economically realistic estimates with reference to an application for which the standard two-stage approach fails to provide plausible insights and would lead to a MFAPM rejection, in a statistical sense. For instance, a two-step F-MB approach leads all the 23 test portfolios to display large, systematic and persistent mis-pricing during our sample period. On the contrary, in the Bayesian case, the values of the posterior medians of the same parameters as well as their signs are sensible and often indicate the absence of large mis-pricings. Third, the F-MB approach shows that idiosyncratic risk is large for most portfolios investigated and highly unstable; in our Bayesian model, when all the uncertainty is taken into account, there is no longer strong evidence of trends in idiosyncratic risk, even though plots for the individual portfolios show some evidence of a peak in the early 2000s and some sign of growth trend towards the end of our sample, consistently with earlier literature (see e.g., Campbell, Lettau, Malkiel and Xu, 2001).

The remainder of the paper is organized as follows. Section 2 outlines the theoretical MFAPM and how we construct factor mimicking portfolios. Section 3 introduces the dynamic Bayesian model with latent stochastic breaks and variances. This section also presents a few competing, restricted versions.

the Bayesian analysis starts and relies on a principal components estimation step. Geweke and Zhou propose a single-step approach but their analysis rules out any instability in betas as well as idiosyncratic risk and does not focus on pre-determined macroeconomic factors.

[4] A few recent papers have used similar time-varying beta multi-factor models with stochastic volatility in a range of applications and provided appropriate Bayesian estimation algorithms. For instance, Aguilar and West (2000) introduce a Bayesian dynamic latent factor model to investigate its portfolio implications; Lopes and Carvalho (2007) have generalized this model to account for breaks in the stochastic volatility process; in a similar context, Carvalho, Lopes and Aguilar (2011) show how latent factors may be combined with observable ones.

Section 4 describes the data and reports the main empirical results. Section 5 performs additional variance decomposition tests. Section 6 concludes.

2. The Pricing Framework

Our empirical work is based on model from the multi-factor linear class introduced by Ferson and Harvey (1991). Multi-factor asset pricing models (MFAPMs) posit a linear relationship between asset returns and a set of (macroeconomic, systematic) factors that are assumed to capture business cycle effects on beliefs and/or preferences (as summarized by a SDF with time-varying properties, see e.g., Cochrane, 2005) and hence on risk premia.[5] If we call the process for the risk factors $F_{j,t}$ $(j = 1, ..., K)$ and $r_{i,t}$ the period *excess* return on asset or portfolio $i = 1, ..., N$, computed as $r_{i,t} \equiv [(P_{i,t} - P_{i,t-1} + D_{i,t})/P_{i,t-1}] - r_t^f$ 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_t^f the one-period interest rate, a typical MFAPM can be written as:

$$r_{i,t} = \beta_{i0,t} + \sum_{j=1}^{K} \beta_{ij,t} F_{j,t} + \epsilon_{i,t} \qquad \epsilon_{i,t} \sim (0, \sigma_{i,t}^2), \tag{1}$$

where $E[\epsilon_{i,t}] = E[\epsilon_{i,t} F_{j,t}] = 0$ for all $i = 1, ..., N$ and $j = 1, ..., K$. The time-varying process for idiosyncratic risk, $\sigma_{i,t}^2$, is left unspecified by asset pricing theory and can be thought as one of the standard frameworks popular in empirical finance, such as a simple GARCH(1,1) or a stochastic volatility model, as it occurs in our paper. The advantage of MFAPMs such as (1) consists of the fact that a number of systematic risk factors $K << N$ may efficiently capture relatively large portions of the variability in the cross-section of returns. Importantly, even though the notation $\beta_{ij,t}$ emphasizes that the factor loadings are allowed to be time-varying, such patterns of time variation are in general left unspecified. Finally, the $\beta_{i0,t}$ coefficients are often interpreted as abnormal returns on asset i "left on the table" after all risks ($F_{j,t}$, $j = 1, ..., K$) and risk exposures ($\beta_{ij,t}$, $j = 1, ..., K$) have been taken into account.

In the conditional version of Ross' (1976) APT (in the absence of arbitrage) or in Merton's (1973) equilibrium intertemporal CAPM (ICAPM), the expected excess return on asset i over the interval $[t-1, t]$ (i.e., the risk premium on asset i), $E_{t-1}[r_{i,t}]$, may then be related by an arbitrage argument to its "betas" (i.e., factor loadings measuring the exposure of asset i to each of the priced, systematic risk factors) and the associated unit risk premia, the $\lambda_{j,t}$s:[6]

[5] The macroeconomic factors with general effects on the SDF are typically represented by the market portfolio (i.e., aggregate wealth) returns, the default spread on corporate bond yields, the term spread incorporated in the riskless (Treasury) yield curve, and changes in the rate of growth of industrial production (see e.g., Chen, Roll and Ross, 1986).

[6] Technically, (2) does not derive from (1) by simply taking conditional expectations. It requires instead assumptions concerning the law of one price, the exact vs. approximate nature of the factor structure, and in the latter case some delicate limiting arguments. Cochrane (2005, ch. 9) provides an introduction to the derivation of MFAPMs starting from a SDF. Under alternative conditions, (1) may also simply hold asymptotically, as $N \to \infty$.

$$E_{t-1}[r_{i,t}] \equiv E[r_{i,t}|\mathbf{Z}_{t-1}] \simeq \lambda_0(\mathbf{Z}_{t-1}) + \sum_{j=1}^{K}\beta_{ij,t|t-1}\lambda_j(\mathbf{Z}_{t-1}) = \lambda_{0,t} + \sum_{j=1}^{K}\beta_{ij,t|t-1}\lambda_{j,t}. \qquad (2)$$

Here both the betas and the risk premia are conditional on the information publicly available at time t, here summarized by the $M \times 1$ vector of "instruments" $\mathbf{Z}_t$ that capture any effects of the state of the economy on unit risk premia (see e.g., Bossaerts and Green, 1989). The framework in (1)-(2) just describes a general conditional pricing framework that is known to hold under a variety of alternative assumptions. However, a variety of methodologies have been proposed to perform three related tasks which affect the empirical performance of (1)-(2):

(i) how many factors ought to be selected, i.e., picking a value for K;

(ii) given K, ranking competing sets of factors;

(iii) estimating the factor loadings $\{\beta_{ij,t}\}$ (over time and for each possible pair i, j) and the risk premia λ_{jt} (over time and for each possible j).

These tasks are logically distinct from the formulation of the framework and—albeit their implementation affects our ability to learn about the fundamental mechanism pricing assets—they have an exquisite statistical nature. In this paper we align ourselves to a number of papers in the empirical finance literature (see e.g., McElroy and Burmeister, 1988; Chen, Roll and Ross, 1986) as far (i)-(ii) are concerned—which means that we pre-select both K and which specific macroeconomic risk factors ought to be considered in the light of the existing literature—and provide an alternative, arguably more flexible econometric approach to accomplish task (iii).

2.1. *The Standard Two-Stage Approach*

The standard approach is the classical, two-stage procedure *à* la Fama and MacBeth (1973) also used by Ferson and Harvey (1991) and popular in empirical finance: In the first step, for each of the assets, the factor betas in (1) are estimated via a simple rolling window OLS. That is, for month t, (1) is estimated using the previous sixty months in order to obtain estimates for the betas, $\hat{\beta}^{60}_{ij,t}$. This time-series regression is updated each month. The choice of a 60-month rolling window scheme is typical of the literature. To favor comparability between our Bayesian implementation with stochastic volatility and the standard two-step approach, all the results in this paper are obtained under the assumption that idiosyncratic variance, $\sigma^2_{i,t}$, follows a standard GARCH(1,1) process, $\sigma^2_{i,t} = \gamma_{i0,t} + \gamma_{i1,t}\epsilon^2_{i,t-1} + \gamma_{i2,t}\sigma^2_{i,t-1}$.[7]

[7] As a result, estimation of both the multi-factor model for the conditional mean and of the variance parameters is performed using quasi-maximum likelihood. Ferson and Harvey (1991) have explored a range of alternative beta estimation techniques, including conditional betas estimated from regressions on past information variables, sixty-month rolling betas regressed on past information variables, and ARCH-style conditional betas, but their results are unaffected by selecting simple, Fama-MacBeth style 5-year rolling OLS regression betas. Guidolin et al. (2013) document that the specifics of the conditional variance model hardly affects the results from the classical, rolling window Fama-MacBeth implementation.

In the second stage, the equilibrium restriction (2) is estimated for each of the periods in our sample a cross-sectional regression using ex-post realized excess returns:

$$r_{i,t} = \lambda_{0,t} + \sum_{j=1}^{K} \lambda_{j,t}\hat{\beta}^{60}_{ij,t} + \zeta_{i,t} \qquad i = 1, ..., N, \quad t = 61, ..., T. \tag{3}$$

Clearly, this T cross-sectional regressions simply implement (2) in a nonparametric fashion, in the sense that any resulting time variation in the $\lambda_{0,t}$ and $\lambda_{j,t}$ coefficients fails to be explicitly and parametrically related to any of the instruments assumed by the researcher, even though additional projections/regressions remain possible. In (3) $\lambda_{0,t}$ is the zero-beta (abnormal) excess return and the $\lambda_{j,t}$s are proxies for the factor risk premiums on each month, $j = 1, ..., K$.[8] Notice that $\lambda_{0,t}$ should equal zero $\forall t$ if the model is correctly specified, because in the absence of arbitrage all zero-beta assets should command a rate of return that equals the short-term rate. Tests of multi-factor models evaluate the importance of the economic risk variables by evaluating whether their risk premiums are priced or whether, on average, the (second-stage, estimated) coefficients $\hat{\lambda}_{j,t}$ are significantly different from zero.

Although widely used, the two-stage Fama-MacBeth (henceforth F-MB) approach has a number of statistical drawbacks. Petersen (2009) discusses these problems in detail and here we limit ourselves to a brief summary, useful to create a contrast with our methodology presented in Section 3. First, the second stage multivariate regression used to test for the equilibrium restriction (2) suffers from obvious generated regressor (error-in-measurement) problems as the estimated first-stage, rolling window beta estimates $\hat{\beta}^{60}_{ij,t-1}$ are used as regressors on the right-hand side. For instance, Ang and Chen (2007) have stressed that when the cross-sectional estimates of the betas $\hat{\beta}^{60}_{ij,t-1}$ co-vary with the underlying but unknown risk premia, (3) may easily yield biased and inconsistent estimates of the risk premia themselves. Unfortunately, this co-variation is extremely likely: for instance, the asset pricing literature generally presumes that during business cycle downturns both the quantity of risk (the size of the betas) and the unit risk prices would increase, simply because recessions are characterized by higher systematic uncertainty as well as by lower "risk appetite" (for instance, in a Campbell and Cochrane's, 1999, habit-formation model). Second, for instance as emphasized by Jostova and Philipov (2005) with reference to a single-factor conditional CAPM, when parameters in linear models are estimated from the data, their uncertainties should be taken into account. Third, the need to perform the estimation of (1)-(2) in two distinct stages that use rolling windows to capture parameter instability is not only *ad hoc* but also inefficient because the lack of more specific parametric forms makes testing for time-variation very hard and dependent on hard-to-justify choices of the rolling window length, the updating rules applied to select whether constant or decaying weights should be applied, etc. (see Maheu and McCurdy, 2009).

[8]This derives from the fact that if one considers a portfolio κ such that $\hat{\beta}^{60}_{\kappa j,t} = 0$ for all $j \neq \kappa$ and $\hat{\beta}^{60}_{\kappa\kappa,t} = 1$, then $\lambda_{\kappa,t}$ is simply the conditional mean of $r_{\kappa,t} - \lambda_{0,t}$.

2.2. *Traded vs. Non-Traded Factors*

One problem with (1) is the difficulty of interpreting $\beta_{i0,t}$ (often called the "Jensen's alpha") when some of the risk factors are not traded portfolios. In principle, $\beta_{i0,t}$ plays a key role: when $F_{j,t} = 0$ for $j = 1, ..., K$, then (1) simplifies to $r_{i,t} = \beta_{i0,t} + \epsilon_{i,t}$ (with $\epsilon_{i,t} \sim (0, \sigma^2_{i,t})$) and any $E[\beta_{i0,t}] \neq 0$ would imply that in the absence of any priced risk factors, the excess return on asset/portfolio i is not zero, which represents a violation of standard economic principles (under the assumption of correct model specification and of a valid implementation/estimation strategy). In this sense, any $E[\beta_{i0,t}] \neq 0$ is referred to as an "abnormal" (average) return. However, although analyses that use (1) to decompose realized excess returns may still be implemented, unless all the factors are themselves tradable portfolios it is impossible to interpret any non-zero β_{i0} as an abnormal return (see Gungor and Luger, 2013). A factor is tradeable if its realizations may be closely replicated ("mimicked", with a high coefficient of determination) by linear combinations (portfolios) of the test assets employed in the analysis. Unless all factors are replicated and replaced by the returns on traded portfolios, there may be a considerable difference between the theoretical alphas from an estimated model, and the actual alpha that an investor may harvest from by trading assets on the basis of a MFAPM.

To eliminate such a possibility, we follow the literature (see e.g., Lamont, 2001) and proceed as follows. When an economic risk factor is already measured in the form of a return (e.g., this is the case of the U.S. market portfolio, real T-bill rates, the liquidity and bond risk factors, term structure spreads, and default spread variables), we directly use the associated returns as a mimicking portfolio. Shanken (1992) has argued that this approach delivers the most efficient estimates of the risk premiums. When a factor is not itself an (excess) return (e.g., this is the case of macroeconomic variables such as industrial production growth, unexpected inflation, and real consumption growth), we construct the corresponding $K' \leq K$ mimicking portfolios by projecting the non-traded factors onto the space of excess returns of base assets and a set of control (predictive) variables ($j = 1, ..., K'$):

$$F_{j,t} = a_j + \mathbf{b}'_j \mathbf{x}_t + \mathbf{c}'_j \mathbf{z}_{t-1} + \boldsymbol{\varepsilon}_{j,t} \qquad \boldsymbol{\varepsilon}_{j,t} \text{ IID } (0, 1), \tag{4}$$

where $\mathbf{x}_t$ is a vector of excess returns on the base assets (in this case, all defined to be zero investment portfolios) and $\mathbf{z}_{t-1}$ denotes a vector of instruments that have the ability to predict returns. The resulting returns on the ith factor mimicking portfolio (FMP henceforth) are then defined as $FMP_{j,t} = \hat{a}_j + \hat{\mathbf{b}}'_j \mathbf{x}_t$ and collect the fitted component of a factor that is unpredictable on the basis of past information and that at the same time may be replicated by trading base assets using weights estimated by $\hat{\mathbf{b}}_j$. Note that the coefficients a_j and $\mathbf{b}_j$ do not need to add up to one because the base assets are zero-investment portfolios (see Lamont, 2001). The base assets include six equity zero net investment portfolios with different book-to-market and size characteristics as well as the returns on long-term government bonds minus the returns on the short term government bonds and the return on long-term corporate bonds minus the return on long-term government bonds. We choose these assets for their well known ability to span large "portions" of the return space. The set of instruments includes the lagged yield spread of long-term Treasury bonds minus the T-bill yield, the lagged yield spread of long-term corporate bonds minus the yield on long-term government bonds, and the lagged real short-term bill rate.

3. **A Bayesian State-Space Approach**

Our discussion of the standard F-MB two-step procedure implies that we need to: (1) avoid using estimates of the first-stage betas as if these were observed variables; (2) fully account for parameter uncertainty; and (3) make an effort to produce a sensible model of parametric instability—here in the form of structural breaks—to reflect the commonly perceived (and tested) fact that both the relationship between excess returns and factors, namely risk exposures ($\beta_{ij,t}$), the risk premia (λ_j, for $i = 1, ..., N$ and $j = 1, ..., K$), and possibly also residual idiosyncratic variances ($\sigma^2_{i,t}$) stochastically change over time, as in Ferson and Harvey (1991). We therefore develop a new Bayesian estimation approach in which:

- The measurement error due to the stochastic nature of the betas is avoided following McCulloch and Rossi (1991) and Geweke and Zhou (1996), by characterizing the joint posterior of risk exposures and premia such that both states and parameters are jointly estimated in a single step.
- Parameter uncertainty is fully addressed by using Bayesian techniques that integrate the joint posterior to find the joint predictive density of the variables of interest.
- Model instability is captured by introducing stochastic breaks in the dynamics of the factor loadings as well as of idiosyncratic volatility.

Specifically, we characterize the relationship between excess returns and factors and the time-varying dynamics in factor loadings and idiosyncratic volatility in a state-space form where the observation equation is the standard linear factor model (1)

$$r_{i,t} = \beta_{i0,t} + \sum_{j=1}^{K} \beta_{ij,t} F_{j,t} + \sigma_{it} \epsilon_{i,t} \tag{5}$$

where $\boldsymbol{\epsilon}_t \equiv [\epsilon_{1,t}, \epsilon_{2,t}, ..., \epsilon_{N,t}]' \sim N(0, \mathbf{I}_N)$ and $E[\epsilon_{i,t}] = E[\epsilon_{i,t} F_{j,t}] = 0$ for all $i = 1, ..., N$ and $j = 1, ..., K$. The time varying parameters $\beta_{ij,t}$ and σ_{it} are described by the state equations

$$\beta_{ij,t} = \beta_{ij,t-1} + \kappa_{ij,t} \eta_{ij,t} \qquad j = 0, ..., K, \tag{6}$$

$$\ln(\sigma^2_{i,t}) = \ln(\sigma^2_{i,t-1}) + \kappa_{i\nu,t} \upsilon_{i,t} \qquad i = 1, ..., N, \tag{7}$$

where $\boldsymbol{\eta}_{i,t} \equiv [\eta_{0,t}, \eta_{1,t}, ..., \eta_{K,t}, \upsilon_{i,t}]' \sim N(0, \mathbf{Q}_i)$ with $\mathbf{Q}_i = diag(q_{i0}^2, q_{i1}^2, ..., q_{iK}^2, q_{iv}^2)$. Stochastic variations (breaks) in the level of both the beta coefficients and of the idiosyncratic variance σ_{it}^2 are introduced and modelled through a mixture innovation approach as in Chan and Maheu (2002), Ravazzolo, Paap, van Dijk and Franses (2007) and Giordani and Kohn (2008). The latent binary random variables $\kappa_{1ij,t}$ and $\kappa_{2i,t}$ are used to capture the presence of random shifts in betas and/or idiosyncratic variance (see Mitchell and Beauchamp, 1988; George and McCulloch, 1993; Miazhynskaia, Frühwirth-Schnatter, and Dorffner, 2006). The random variable $\kappa_{1ij,t}$ takes then a value equal to one if a structural break for the jth factor in the equation for the ith asset at time t takes place. We assume that the structural breaks are independent of each another (i.e., across assets as well as factors) and over time, with:[9]

$$\Pr\left[\kappa_{ij,t} = 1\right] = \pi_{ij} \qquad \Pr\left[\kappa_{i\nu,t} = 1\right] = \pi_{i\nu} \qquad i = 1, ..., N \quad j = 0, ..., K \tag{8}$$

This specification is very flexible as it allows for both constant and time-varying parameters. When $\kappa_{ij,\tau} = \kappa_{i\nu,\tau} = 0$ for some $t = \tau$, then (6) reduces to (1) when the factor loadings and the quantity of idiosyncratic risk are assumed to be constant, as $\beta_{ij,\tau} = \beta_{ij,\tau-1}$ and $\ln \sigma_{i,\tau}^2 = \ln \sigma_{i,\tau-1}^2$. However, when $\kappa_{1ij,\tau} = 1$ and/or $\kappa_{2i,\tau} = 1$, a break hits either a beta or idiosyncratic variance or both, and instability is then captured by the random walk dynamics $\beta_{ij,\tau} = \beta_{ij,\tau-1} + \eta_{ij,\tau}$ and $\ln(\sigma_{i,\tau}^2) = \ln(\sigma_{i,\tau-1}^2) + \upsilon_{i,t}$ (or $\sigma_{i,\tau}^2 = \sigma_{i,\tau-1}^2 \exp(\upsilon_{i,\tau})$). The flexibility of the specification in (6) stems from the fact that risk exposures, $\beta_{ij,t}$, and idiosyncratic risks, $\sigma_{i,t}^2$, are allowed to change on every time period, but they are not imposed to be changing at every point in time. In our view, this helps to side-step the difficult (if not impossible) task of persuading a Reader that the assumed dynamics represents the "right" kind: given our uninformative priors, if the data need frequent breaks in betas of a small size, the posterior of the corresponding parameters will provide indications in this direction; similarly, if the data need a (set of independent) stochastic volatility process(es) with frequent shifts in idiosyncratic variance, posterior estimates will give appropriate indications, etc.

Note that because when a break affects the betas and/or the variances, the random shift is measured by variables collected in $\boldsymbol{\eta}_{i,t}$, we can interpret $\mathbf{Q}_i$ as the "size" of the break: a large q_{ij}^2 means for instance that whenever $\beta_{ij,t}$ is hit by a break, i.e. $\kappa_{ij,t} = 1$, such a shift is more likely to be large (in absolute value). This process for factor loadings and idiosyncratic residual risk is different from frameworks typical of the time-varying parameter literature in which factor loadings are assumed to vary continuously (i.e., in every period) and usually according to simplistic AR(1) structures with high persistence and small variance for the shocks, such as $\beta_{ij,t} = \phi_{1ij}\beta_{ij,t-1} + \eta_{ij,t}$ and $\ln(\sigma_{i,t}^2) = \phi_{2i} \ln(\sigma_{i,t-1}^2) + \upsilon_{i,t}$ with ϕ_{1ij} and

[9] The independence across breaks is consistent with the spirit of a factor model and may not be necessarily restrictive. Indeed, the comovements among asset returns should be driven by the factor structure regardless of the nature of the structural breaks. However our approach is flexible because "it lets the data speak" about whether breaks across assets i and j are contemporaneous or not.

ϕ_{2j} close to but less than one.[10]

Following a logic similar to (3) but applied to our B-TVB-SV framework, the cross-sectional restrictions in (2) are characterized through the multivariate linear model

$$r_{i,t} = \lambda_{0,t} + \sum_{j=1}^{K} \lambda_{j,t}\beta_{ij,t|t-1} + e_{i,t} \qquad i = 1, ..., N, \tag{9}$$

where $e_{i,t} \sim N(0, \psi^2)$ and $\beta_{ij,t|t-1}$ measures the *expected* time t sensitivity of asset i to factor j, based on all information available up to time $t-1$. Because we adopt a single-step Bayesian estimation strategy, and the unknown betas and risk premia are multiply each other, in practice (9) imposes a set of non-linear restrictions in estimation. Note that under the assumption of correct specification of the asset pricing model, $\lambda_{0,t} = 0$ or, at least, the average over time of the $\lambda_{0,t}$s ought to be zero, implying that only the assumed risk factors are explaining the risk premia on the different assets and portfolios. In our setting $\beta_{ij,t|t-1}$ represents a draw from the predictive distribution of the state dynamics in (6), which is obtained by integrating out both the probability of recording a structural break and the uncertainty about the size of the break itself. Ferson and Harvey (1991) have emphasized the importance that in the implementation of factor models the time t excess return on asset i should be determined by investors with reference only to information available up to time $t-1$.[11] The Bayesian paradigm allows us to go one step further to properly capture the forward looking nature of the asset pricing model. As in Geweke and Zhou (1996), the risk premia $\boldsymbol{\lambda}_t \equiv (\lambda_{0,t}, \lambda_{1,t}, ..., \lambda_{K,t})'$ are estimated jointly with the factor loadings $\mathbf{B}_t \equiv \left\{\beta_{ij,t}\right\}_{i=1}^{N}{}_{j=0}^{K}$, the (log of the) idiosyncratic variances $\boldsymbol{\sigma}_t^2 \equiv \left(\sigma_{1t}^2, \sigma_{2t}^2, ..., \sigma_{Nt}^2\right)'$, as well as the other parameters $\boldsymbol{\Theta} = \{\boldsymbol{\theta}_i\}_{i=1}^{N}$, with $\boldsymbol{\theta}_i \equiv (\mathbf{q}_i^2, \boldsymbol{\pi}_i)'$, where $\mathbf{q}_i^2 \equiv \left(q_{i0}^2, q_{i1}^2, ..., q_{iK}^2, q_{iv}^2\right)'$ is the vector of conditional variances of the factor loadings and the idiosyncratic risks and $\boldsymbol{\pi}_i \equiv (\pi_{i0}, ..., \pi_{iK}, \pi_{i\nu})'$ is the vector of structural break probabilities for the ith asset. The time variation in the risk premia is inherited by the dynamics in portfolio sensitivities, $\{\mathbf{B}_t\}_{t=1}^{T}$. Therefore, even though the dynamics of $\boldsymbol{\lambda}_t$ is not explicitly specified in our model, the instability of the betas is by construction reflected in the risk premia as well. Appendix A provides additional details on the estimation algorithm.

3.1. *Special Cases*

The model presented in (5)-(8) is the most general specification we consider in this paper. We will call this model B-TVB-SV specification indicating that we consider a Bayesian (B), Time-Varying Betas (TVB) and Stochastic Volatility (SV) framework. Here the words time-varying and stochastic for the betas and

[10] The likelihood tends to be not well-behaved when ϕ_{1ij} and ϕ_{2i} are close to one and their estimation might be difficult, see the discussion and examples in De Pooter, Ravazzolo, Segers and van Dijk (2008).

[11] The parameters $\mathbf{Q}_i$ and probabilities $\Pr(\kappa_{1ij,\tau} = 1)$ and $\Pr(\kappa_{2i,\tau} = 1)$ are however estimated over the full sample period. It is possible to recursively repeat the estimation over several vintages of data and produce out-of-sample forecasts of $\beta_{ij,t|t-1}$, but the computational cost in our application with 40 years data and 23 portfolios would be very high.

the volatilities are synonymous of structural breaks in both the risk exposures and the idiosyncratic risks. Of course, this B-TVP-SV model is richly parameterized and it cannot be ruled out that issues related to over-parameterization may arise. Moreover, many of our fine economic conclusions might be driven by details of the parameterization of the change point process in (5)-(8). Therefore, for comparative purposes, we consider a number of alternative restrictions on the dynamics of the state equation:

1. $\kappa_{i\nu,t} = 0 \ \forall i,\ t$, i.e. a constant idiosyncratic volatility model:

$$\begin{aligned} r_{i,t} &= \beta_{i0,t} + \sum_{j=1}^{K} \beta_{ij,t} F_{j,t} + \sigma_i \epsilon_{i,t} \qquad i = 1, ..., N, \\ \beta_{ij,t} &= \beta_{ij,t-1} + \kappa_{ij,t} \eta_{ij,t} \qquad j = 0, ..., K, \end{aligned} \tag{10}$$

   under the same distributional assumption as (5)-(9). We will call this model a Bayesian homoskedastic time-varying betas model, i.e. B-TVB.[12]

2. $\kappa_{ij,t} = 1 \ \forall i,\ j,\ t$ and $\kappa_{i\nu,t} = 1 \ \forall i, t$, i.e., time-varying parameters (TVP) according to random walk specifications (see e.g., Koop and Potter, 2007; West and Harrison, 1997). In practice, we also implement a Bayesian TVP model (B-TVP) also considered by Jostova and Philipov (2005):

$$\begin{aligned} r_{i,t} &= \beta_{i0,t} + \sum_{j=1}^{K} \beta_{ij,t} F_{j,t} + \sigma_{it} \epsilon_{i,t} \\ \beta_{ij,t} &= \beta_{ij,t-1} + \eta_{ij,t} \qquad j = 0, ..., K, \\ \ln(\sigma_{i,t}^2) &= \ln(\sigma_{i,t-1}^2) + \upsilon_{i,t} \qquad i = 1, ..., N. \end{aligned} \tag{11}$$

   B-TVP assumes a unit probability of breaks (even though this are of a small size) in the dynamics of the states $\beta_{ij,t}$ and $\sigma_{i,t}^2$ at each point in time. This is indeed a fairly strict assumption which is not necessarily supported by the data. Even though we name the model B-TVP, it features SV.

Of course, the constant volatility B-TVB specification is used to highlight the effects of instabilities in residual variances. The B-TVP specification is used as a competing specification in order to show the benefit of considering the more parsimonious, occasional large breaks in (6)-(8) as opposed to small, frequent (continuous) breaks (see Giordani and Villani, 2010, for a related discussion).

3.2. *Prior Specification*

We estimate (6) using a Bayesian approach that allows us to incorporate parameter uncertainty when estimating both the beta exposures and the equilibrium risk premia. For the Bayesian algorithm illustrated

[12] Trivially, the symmetric case of $\kappa_{ij,t} = \kappa_{iv,t} = 0 \ \forall t$ implies that $\beta_{ij,t} = \beta_{ij,t-1} = \beta_{ij}$ and $\ln(\sigma_{i,t}^2) = \ln(\sigma_{i,t-1}^2) = \ln(\sigma_i^2)$ and consists of the two-step Fama-MacBeth model with constant betas and idiosyncratic variances.

in Appendix A to work, we need to specify the prior distributions for each of the model parameters. The parameters of the model (5)-(9) are $\mathbf{\Theta} \equiv \{\boldsymbol{\theta}_i\}_{i=1}^N$ with $\boldsymbol{\theta}_i \equiv (\mathbf{q}_i^2, \boldsymbol{\pi}_i)$, plus the risk premia $\boldsymbol{\lambda}_t$ which are estimated at each time t conditional on the factor exposures sensitivities according to (9). We choose a conjugate prior structure to keep the numerical analysis as simple as possible. As far as the structural break probabilities are concerned, we assume a set of simple Beta distributions:

$$\pi_{ij} \sim Beta(a_{ij}, b_{ij}) \qquad \pi_{i\nu} \sim Beta(a_{i\nu}, b_{i\nu}) \quad \text{for} \quad i = 1, ...N, \quad j = 1, ...K. \tag{12}$$

The parameters a_{ij}, b_{ij} and $a_{i\nu}, b_{i\nu}$ represent the shape hyperparameters and can be set according to our prior beliefs about the occurrence of structural breaks in $\beta_{ij,t}$ and $\ln(\sigma_{i,t}^2)$, respectively.[13]

For the conditional variance parameters $\mathbf{q}_i^2$, which reflect our prior beliefs about the size of the structural breaks we assume an inverted Gamma prior,

$$q_{ij}^2 \sim IG(\gamma_{ij}, \delta_{ij}) \qquad q_{i\nu}^2 \sim IG(\gamma_{i\nu}, \delta_{i\nu}) \quad \text{for} \quad i = 1, ...N, \quad j = 1, ...K \tag{13}$$

where $\gamma_{ij} > 0$, $\gamma_{i\nu} > 0$ and $\delta_{ij} > 2$, $\delta_{i\nu} > 2$ are the scale and degrees of freedom parameters, respectively, for the factor loadings and the (log-) variances.[14] Finally, the prior distribution for the risk premia $\boldsymbol{\lambda}_t$ is a characterized as a standard multivariate normal distribution with independent priors:

$$\boldsymbol{\lambda}_t \sim N_K(\underline{\lambda}, \underline{V}) \qquad \psi^2 \sim IG(\psi_0, \Psi_0) \quad \text{for} \quad t = 1, ...T. \tag{14}$$

The parameters $\underline{\lambda}$ and $\underline{V}$ represent the $K \times 1$ location vector and the $K \times K$ scale matrix for the K-dimensional multivariate normal distribution; ψ_0 and Ψ_0 are the scale and degrees of freedom of the conditional variance ψ^2 parameters, respectively, in (9). Because these priors are independent of one another, the density of the joint prior distribution $p(\mathbf{\Theta})$ is given by the product of the prior specifications (12)-(14). The choice of the values for the hyperparameters of the priors is discussed in Appendix A.

3.3. *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 $\beta_{ij,t}$, σ_{it}^2 and $\kappa_{ij,t}$, $\kappa_{i\nu,t}$ for each of the $i = 1, ..., N$ assets, each of the $j = 1, ..., K$ factors and at each time $t = 1, ..., T$, are simulated alongside the model parameters $\boldsymbol{\theta}_i$ and the equilibrium risk premia $\boldsymbol{\lambda}_t$. However, to apply the Gibbs sampler we need to

[13]Under a Beta distribution, the unconditional expected prior probability of a structural break for the ith asset beta relative to the jth factor is defined as $a_{ij}/(a_{ij}+b_{ij})$ while in the case of idiosyncratic variance, this is equal to $a_{i\nu}/(a_{i\nu}+b_{i\nu})$.

[14]Under an Inverted Gamma prior, the expected size of a break for, say, the exposure of ith asset to the jth factor is $\gamma_{ij}/(\delta_{ij} - 2)$ for $\delta_{ij} > 2$.

write down the complete likelihood function, namely, the joint density of the data and the state variables. Defining $\boldsymbol{\theta} \equiv \{\boldsymbol{\theta}_i\}_{i=1}^{N}$, $\mathbf{B}_t \equiv \{\boldsymbol{\beta}_{it}\}_{i=1}^{N}$, $\mathbf{B} \equiv \{\mathbf{B}_t\}_{t=1}^{T}$, $\mathbf{R} \equiv \{r_{it}\}_{i=1}^{T}{}_{t=1}^{N}$, $\mathbf{F} \equiv \{\boldsymbol{F}_t\}_{t=1}^{T}$, $\boldsymbol{\lambda} \equiv \{\boldsymbol{\lambda}_t\}_{t=1}^{T}$, $\mathcal{K}_\beta \equiv \{\kappa_{ij,t}\}_{j=1}^{K}{}_{i=1}^{N}{}_{t=1}^{T}$, $\mathcal{K}_\sigma \equiv \{\kappa_{iv,t}\}_{i=1}^{N}{}_{t=1}^{T}$, $\boldsymbol{\Sigma} = \{\boldsymbol{\sigma}_{it}^2\}_{i=1}^{N}{}_{t=1}^{T}$, the likelihood function is

$$p(\mathbf{R}, \mathbf{B}, \mathcal{K}, \boldsymbol{\Sigma}, \boldsymbol{\lambda}|\boldsymbol{\theta}, \mathbf{F}) = \prod_{t=1}^{T} \left\{ \prod_{i=1}^{N} p(r_{it}|\mathbf{F}_t, \boldsymbol{\beta}_{it}, \sigma_{it}^2) p(\sigma_{it}^2|\sigma_{it-1}^2, \kappa_{i\nu,t}, q_{i\nu}^2) \pi_{i\nu}^{\kappa_{i\nu t}} (1-\pi_{i\nu})^{1-\kappa_{i\nu t}} \times \right. \tag{15}$$

$$\left. \times \left[\prod_{j=0}^{K} p(\beta_{ij,t}|\beta_{ij,t-1}, \kappa_{ij,t}, q_{ij}^2) \times \pi_{ij}^{\kappa_{ijt}} (1-\pi_{ij})^{1-\kappa_{ijt}} \right] p\left(\boldsymbol{\lambda}_t, \psi^2|\mathbf{B}_t, \mathbf{R}_t\right) \right\},$$

where $\mathcal{K} \equiv (\mathcal{K}_\beta, \mathcal{K}_\sigma)$ and $\boldsymbol{F}_t = (F_{1,t}, F_{2,t}, ..., F_{K,t})'$. Combining the prior specifications (12)-(14) with the complete likelihood, we obtain the posterior density $p(\boldsymbol{\theta}, \mathbf{B}, \mathcal{K}, \boldsymbol{\Sigma}, \boldsymbol{\lambda}|\mathbf{R}, \mathbf{F}) \propto p(\boldsymbol{\theta}) p(\mathbf{R}, \mathbf{B}, \mathcal{K}, \boldsymbol{\Sigma}, \boldsymbol{\lambda}|\boldsymbol{\theta}, \mathbf{F})$. Our Gibbs sampler is a combination of the Forward Filtering Backward Sampling of Carter and Kohn (1994) and Kim, Shepard, and Chib (1998), and the efficient sampling algorithm for the random breaks proposed in Gerlach, Carter, and Kohn (2000). At each iteration of the sampler we sequentially cycle through the following steps:

1. Draw $\mathcal{K}_\beta$ conditional on $\boldsymbol{\Sigma}, \mathcal{K}_\sigma, \boldsymbol{\theta}$, $\mathbf{R}$ and $\mathbf{F}$.
2. Draw $\mathbf{B}$ conditional on $\boldsymbol{\Sigma}, \mathcal{K}, \boldsymbol{\theta}$, $\mathbf{R}$ and $\mathbf{F}$.
3. Draw $\mathcal{K}_\sigma$ conditional on $\mathbf{B}, \mathcal{K}_\beta, \boldsymbol{\theta}$, $\mathbf{R}$ and $\mathbf{F}$.
4. Draw $\mathbf{R}$ conditional on $\mathbf{B}, \mathcal{K}, \boldsymbol{\theta}$, $\mathbf{R}$ and $\mathbf{F}$.
5. Draw $\boldsymbol{\lambda}$ conditional on $\mathbf{B}, \mathcal{K}, \boldsymbol{\theta}$, $\mathbf{R}$ and $\boldsymbol{\Sigma}$.
6. Draw $\boldsymbol{\theta}$ conditional on $\mathbf{B}, \mathcal{K}$, $\mathbf{R}$ and $\mathbf{F}$.

We use a burn-in period of 1,000 and draw 5,000 observations storing every other of them to simulate the posterior distribution of parameters and latent variables. The resulting autocorrelations of the draws are very low.[15] Appendix A provides additional details.

4. An Empirical Application to the U.S. Cross-Section of Financial Returns

4.1. *Data and Descriptive Statistics*

We consider a typical application in the empirical finance literature based on a large number (23) of monthly time series sampled over the period 1972:01 - 2011:12. The starting date is due to the availability of the complete set of instruments and corporate bond return data. The initial ten years are used to empirically elicit the priors. Our empirical analysis is implemented over the remaining 360 observations,

[15] In 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.

per each of the series. The series belong to two main categories. The first group, "Portfolio Returns", includes stocks, U.S. Treasuries and notes, and corporate bonds, all organized in portfolios to tame the non-diversifiable risk reflected by excess returns. The stocks are publicly traded firms listed on the NYSE, AMEX and Nasdaq (from CRSP) and sorted according to two criteria. First, 10 industry portfolios are obtained by sorting firms according to their four-digit SIC code. Second, 10 additional portfolios are derived by sorting (at the end of every year, and recursively updating this sorting every year) NYSE, AMEX and Nasdaq stocks according to their size, as measured by the aggregate market value of the company's equity. Using industry and size-sorting criteria to form portfolios of stocks to trade-off "spread" and reduction of idiosyncratic risk, is typical in the literature (see e.g., Dittmar, 2002). Moreover, industry- and size-sorting criteria are sufficiently unrelated to make it plausible that industry- and size-sorted equity portfolios may contain non-overlapping information on the underlying factors and risk premia. Data on long- (10-year) and medium-term (5-year) government bond returns are from Ibbotson and available from CRSP. Data on 1-month T-bill, 10-year and 5-year government bond yields and returns are from FREDII at the Federal Reserve Bank of St. Louis and from CRSP. Data on "junk" bond returns are approximated from Moody's (10-to-20 year maturity) Baa average corporate bond yields and converted into return data using Shiller's (1979) approximation formula.

The second group collects macroeconomic risk variables. These factors are used as proxies for the systematic, economy-wide forces potentially priced in asset returns. We employ nine factors: the excess return on a wide, value-weighted market portfolio (r_t^M) that includes all stocks traded on the NYSE, AMEX, and Nasdaq (from CRSP); changes in the default risk premium (def_t) measured as the difference between Baa Moody's *yields* and yields on 10-year government bonds; the change in the term premium ($\Delta term_t$), the difference between 10-year and 1-month Treasury yields; the unexpected inflation rate ($UInfl_t$), computed as the residual of a simple ARMA(1,1) model applied to (seasonally adjusted) CPI inflation rate; the rate of growth of (seasonally adjusted) industrial production (IP_t); the rate of growth of (seasonally adjusted) real personal consumption (PC_t); the 1-month real T-bill return computed as the difference between the 1-month T-bill nominal return and realized CPI inflation rate (not seasonally adjusted); the traded Liquidity factor (Liq_t) from Pastor and Stambaugh (2003); the Bond premium factor ($Bondp_t$) from Cochrane and Piazzesi (2005).[16] Using a relatively large number of pre-selected factors is typical of the literature.[17] Table 1 reports a detailed set of summary statistics.

[16] The traded liquidity factor consists of value-weighted returns on a high-minus-low exposure portfolio on an aggregate liquidity risk factor that sorts stocks on the basis of liquidity risk measures. Næs, Skjeltrop, and Ødegaard (2011) show the existence of strong linkages between stock market liquidity and business cycle-related macroeconomic aggregates. The bond risk premium factor is constructed as the projection of the equally weighted average of one-year excess holding period return on bonds with maturities of two, three, four, and five years on a constant, the one-year yield, and the two- through five-year forward rates. The bond risk factor is the fitted value of this regression. Cochrane and Piazzesi (2005) and Ludivgson and Ng (2009) investigate the relationship beween this factor and macroeconomic aggregates, following the intuition of Harvey (1989) on the links between the term structure and consumption growth.

[17] For instance Mei (1993) uses five factors; Connor and Korajczyk (1988) find there are more than five factors (although factors in excess of five generally do not play an important role, although they are statistically significant); Ludvigson and

4.2. *Time-Varying Betas*

As an initial way to assess the plausibility of our results, Figure 1 reports the average (of posterior medians over time) probabilities over our sample of observing a break in the factor loadings, in addition to the intercept, across two different specifications, namely the B-TVB-SV and the homoskedastic B-TVB, for the 23 test assets/portfolios. Clearly the presence of breaks in the idiosyncratic variance process makes a difference in capturing any instability in portfolio betas. Under the B-TVP-SV model the average probability of observing a break is around 40% for the intercept (labeled as factor 1 in the figure) of all portfolios examined, and ranges from 20% for the credit and term spreads (factors 2 and 8) to almost 40% for the bond factor (factor 9). This shows that infrequent and large breaks in betas (as well as Jensen's alphas) are often isolated by the Gibbs sampling algorithm. Under the B-TVB specification, instead, the degree of instability in the factor loadings dramatically collapses. The average probability of a break in betas is around 5% across all risk for the industry portfolios (portfolios 1-10 in the figure), while for both the size-sorted equity portfolios and bonds, the average break probability over the sample increases to between 20 and 30% across factors.

Figures 2-7 plot a selection of time series medians and 95% Bayesian credibility intervals computed from the posterior densities of the loadings $\beta_{ij,t}$, obtained from the B-TVB-SV model. To save space, we report plots of time series of risk exposures for all the 23 portfolios used in our estimation, but only for five out of nine specific factors: the U.S. market portfolio, the term spread, industrial production growth, the real T-bill, and unexpected inflation. Other, similar plots concerning the remaining risk factors—the credit spread, the real consumption growth, the bond and the liquidity factors —are available upon request even though we summarize their contents and implications below. An overview of the plots immediately reveals that the Bayesian estimates of the loadings for all but the market portfolio and the bond risk factor, imply a time path of the factor loadings that is rather smooth over time. This is a first interesting result: even though (6) formally allows factor exposures to be subject to "jumps" over time, as a result of the realization of $\kappa_{ij,t}$, the resulting posterior densities are actually smoother than what one could retrieve using, say, a naïve rolling window scheme. For instance, this is evident from a comparison of Figures 2 and 3, where in the latter we plot estimated, 5-year rolling window F-MB betas for the 23 test portfolios vs. the market portfolio.[18] Interestingly, this smoothness mimics exactly what many earlier papers have imposed by assuming near unit root processes ($\beta_{ij,t} = \phi_{ij}\beta_{ij,t-1} + u_{ij,t}$) with small variance of the shocks, but is derived endogenously and is data-driven, which means that occasional large jumps in exposures and/or high volatility of the process may be accommodated. Second, with a limited number of exceptions that will noted below, the 95% confidence bands are relatively tight, which means that the betas are estimated with a fairly high level of reliability.

Ng (2009) find evidence in favor of eight latent factors.

[18] To save space, we do not report all plots of the time-varying, 5-year rolling window betas obtained using a classical two-step estimation scheme, as we describe in Appendix A. These plots are available upon request.

In particular, Figure 2, concerning exposures to market risk, collects most of the loadings for which we have evidence that betas are non-zero. All equity portfolios are characterized by positive and reliably estimated betas. This is not the case for the bond portfolios which essentially show zero exposure to the market risk factor. As already mentioned, Figure 3 offers an opportunity to compare the B-TVB-SV estimates with market beta exposures under a the classical F-MB approach described in Appendix A. In Figure 4, concerning the betas vs. term premium shocks, most equity portfolios are significantly exposed to yield curve slope risk, in the sense that their 95% bands do not systematically include zero. In the figure, the betas fluctuate considerably over the sample period and often change sign. Such betas tend to drift down and to be (significantly) negative for low-decile size-sorted portfolios (i.e., small and medium stocks), for high-tech stocks (especially after 1994), and for junk corporate bonds (at least in the 1980s and 1990s); they are instead positive and often significant in the case of energy and health stocks.

The plots of time-varying exposures to real output (industrial production growth) risk in Figure 5 show occasionally larg(er) 95% credibility regions that tend to widen over the sample. However, also in this case, for a large sub-set of portfolios, the corresponding betas are estimated to be negative and significant (nondurables, durables, manufacturing, high-tech, shops, health, and small- and medium-size equity portfolios), while for other portfolios the exposure is positive and significant (energy and utility stocks). Of course, negative exposures to output risk are partially surprising, but because in our model, factors have not been orthogonalized one vs. the others—that will require selecting and imposing a triangular structure that would prove to be "ad hoc"—betas only capture partial effects, after other exposures to business cycle risks are taken into account (see Kramer, 1994). An unreported figure concerning betas vs. the short term real rate shows instead exposures that are small and for which the 95% credibility bands tend to include zero for most of the sample. However, close attention reveals that a number of smooth patterns of fluctuations imply sub-periods in which exposures to real rate risk have also been precisely estimated, for instance a negative exposure in the case of durables, high tech, retail shop, first capitalization decile stocks, and medium-term Treasury notes, in the early and mid-1990s.

Figure 6 shows estimated time-varying exposures to unexpected inflation risks. In the asset pricing literature, the issue of the exposure of asset returns to inflation risks has often been debated. The plots show that even though confidence bands tend to be wider for this factor than for other factors that we have described before, for many portfolios there tends to be still significant evidence of a significantly positive exposure, i.e., of the fact that these assets pay out risk premia to compensate for inflation risks. Even if we limit ourselves to global results that hold throughout our entire sample, this hedging property obtains in the case of durables, high-tech, retail, and of small and medium-capitalization stocks. On the contrary, energy, telecommunication, utilities, and especially all kinds of bonds (including corporate junk), imply negative, significantly estimated exposures throughout the sample.

We have also inspected the remaining four sets of plots concerning the other risk factors (default spread, real consumption growth, liquidity, and bond factors). In the case of default spread and real consumption growth risks, all betas imply low variability and narrow 95% credibility regions, but these also fluctuate steadily around zero for the *all* 23 portfolio investigated. The finding is similar in the case of betas vs. the liquidity factor, although these show a stronger temporal variation, which means that sub-samples can be isolated in which this factor commands precisely estimated exposures. When this occurs, betas are mostly positive (for instance, durable stocks in the late 1990s and high-tech stocks in the early 1990s). Finally, although unstable, the results concerning the exposures to Cochrane and Piazzesi's factor are interesting. As one would expect, all bond portfolios show periods in which they have large and significant loadings on this factor, especially during the 1990s and 2001-2004. Equities show more dispersed posteriors which become large in the 1990s.

Figure 7 reports posterior medians and 95% credibility intervals for the $\beta_{i0,t}$s estimated from the B-TVB-SV model. In an ICAPM interpretation of (1), made possible by the fact that all factors are traded, the time series (of posterior medians for) $\{\beta_{i0,t}\}$ gives indications on mis-pricing. Out of 23 portfolios, in no case the estimated mis-pricing indicators are systematically elevated (in absolute) value. In fact, apart from occasional fluctuations, separate calculations show that the 95% credibility regions include a zero mis-pricing in more than three-quarters of our sample. This is an indication that in its B-TVB-SV implementation, the model (5)- (6) is well-specified in an economic sense, as it does not imply any evidence of a systematic mis-pricing. Of course, in the case of many portfolios, occasional periods in which the posterior of $\beta_{i0,t}$ fails to include a zero mis-pricing can be found. For instance, there is evidence that all bond portfolios implied positive and tight posteriors for the Jensen's alphas between 2000 and 2004; high-tech and telecommunication stocks were all giving large and significant alphas during the early- to mid-1990s. Interestingly, Figure 8 shows that this is not the case when we plot and examine estimates of the $\beta_{i0,t}$ coefficients from a 5-year F-MB implementation. The corresponding Jensen's alphas are strongly fluctuating often reaching extreme levels of $\pm 4\%$ *per month*, i.e., values that are hardly plausible in an asset pricing perspective, and they often appear to be statistically significant in the sense that their 95% confidence bands fail to include a zero mis-pricing. This is of course hardly credible and points more to a failure of the two-pass, rolling window approach than to a misspecification issue with the model in (1).

4.3. *Dynamics in Idiosyncratic Risk*

A growing literature (see e.g., Campbell, Lettau, Malkiel and Xu, 2001) has stressed that the idiosyncratic variance of the excess returns on most test portfolios, σ^2_{it}, has undergone important shifts and/or dynamics over the last two decades. We have inspected the filtered GARCH values of residual variance obtained from the classical, two-pass F-MB method (unreported to save space). The presence of rich dynamics is obvious for all the portfolios. In some cases such dynamics turned out to be hard to interpret. First, idiosyncratic variation tends to be large for most test assets, a sign that the two-pass method provides does not fully explain the 23 time series of excess returns. The ratio between the average of the two-pass GARCH estimates $\overline{\hat{\sigma}_{it}^2}$ and the sample variance of portfolio returns over our sample reveals that for a large fraction (13) of the 23 portfolios investigated, this is close to or in excess of 0.5, with all the bond portfolios with ratios in excess of one (because of the rolling window nature of the $\{\hat{\sigma}_{it}\}$, this ratio may actually exceed one). This means that the non-systematic component of excess returns still explains at least 50% of the total variance. Second, most equity portfolios (in practice, all the industry portfolios and size deciles 1-8) record a peak in idiosyncratic variance between 2000 and 2003. In some cases, the rolling window idiosyncratic variance practically doubles between 1999 and 2004, meaning that the MFAPM loses most of its ability to fit excess returns using risk exposures.

Figure 9 plots posterior medians for σ_{it}^2 estimated from (5)- (6), along with 95% Bayesian intervals. There are evident spikes in idiosyncratic volatility in the early 2000s and weaker signs of a growing trend towards the end of our sample. The financial crisis of 2008-2009 induces a residual risk increase, but this appears to be minor compared to 1999-2001, when the model had temporarily lost its ability to fit the U.S. cross-section. In this respect, the fact that the model is more at trouble with the tech stock bust than with the U.S. subprime and credit crunch crises is intriguing. However, the fact that idiosyncratic is countercyclical was largely expected in the light of the literature (see Campbell et al., 2001). The B-TVB-SV model explains away almost all the variability in excess returns in the case of medium and large cap stocks, and to some extent also government bonds. Spikes in idiosyncratic risk are instead more pronounced for small caps and for a number of industry portfolios, that are explained much less accurately than size-sorted portfolios are. Yet, no clear trend is observed, which is consistent with the more recent evidence reported by Bekaert, Hodrick, and Zhang (2012).

4.4. *Time-Varying Risk Premia*

Table 2 reports summary statistics for posterior estimates of the risk premia $\{\hat{\lambda}_{j,t}\}$ $(j = 1, ..., K)$ from the B-TVP-SV model as well as the B-TVP and the homoskedastic B-TVB frameworks. As a benchmark, the table also shows (frequentist) estimates from the second-pass F-MB approach. In the table, we also report the empirical standard error for the sample mean of each of the $\lambda_{j,t}$s. From the very first panel of the table, it is clear that the classical estimation procedure that non-parametrically tracks time-variation in the parameters using 5-year rolling windows delivers economically weak implications: only two factors were accurately priced in the cross-section (the market and bond factors), but the former with a p-value exceeding the standard 0.05 threshold and the latter with a rather difficult, negative sign; moreover, the time series mean of $\hat{\lambda}_{0,t}$ turns out to be large, positive (0.29% per month), and statistically significant (its p-value is 0.034), which is problematic to our MFAPM because a non-zero average $\lambda_{0,t}$ implies that omitted risk factors with non-zero risk premia must be absorbed by the residual mean.

Therefore the background to our dynamic, state-space results is that a simple, *ad-hoc* rolling window implementation of the MFAPM in (5)- (6) would yield an embarrassing rejection of the model, in spite of the fact that we are employing as many as nine factors, some of them coming with a strong endorsement of cross-sectional explanatory power from the asset pricing literature. Fortunately, a much more comforting picture emerges from the B-TVB-SV model, when a parametric model for unstable risk exposures and idiosyncratic risk is assumed. Here the Bayesian design gives evidence of precisely estimated market, liquidity, and macroeconomic (as capture by IP growth shocks) risk premia, with the correct, positive signs (0.339, 0.317, and 0.002 percent per month/unit of risk, respectively). Also the unexpected inflation risk premium is precisely estimated but with a negative sign, similarly to Chen, Roll, and Ross (1986), Ferson and Harvey (1991), and Lamont (2001). Importantly, the average of the posteriors for $\lambda_{0,t}$ reveals that the intercept is not significantly different from zero.

All in all, these results illustrate the fact that while in a naive F-MB implementation all one gets is evidence that a standard multi-factor model—both in terms of its structure and for what concerns the factor it includes—is rejected with reference to a wide but typical set of U.S. financial asset portfolios, such finding is replaced by reassuring evidence that not only the market portfolio (as typical of text-book CAPM) but also a number of macroeconomic factors carry precisely estimated and economically meaningful risk prices. Such empirical findings are less comforting when we impose restrictions on the B-TVP-SV model. In both cases, the average of the posteriors for $\lambda_{0,t}$ has a mean that is significantly positive, with p-values below 0.05. While in the B-TVP case at least 3 of the 4 factors that commanded positive and significant risk premia in B-TVP-SV set up, in the homoskedastic B-TVB model only market risks appear to be barely priced in the U.S. cross-section of stocks and bonds. This is indicative of the restrictions imposed by the B-TVP and the homoskedastic B-TVB models being rejected, an aspect that shall be investigated in more detail in Section 4.5.

4.5. *Discriminating Among Models: Marginal Likelihood Evidence*

Following McCulloch and Rossi (1991), we use the marginal likelihood of different models to perform a comparison able to take into account their overall (in-sample) statistical performance, and not only their asset pricing plausibility as in Sections 4.2-4.4. The marginal likelihood of a model is known to take into account both the uncertainty about the size and the presence of structural breaks and the uncertainty concerning the parameters in (5)- (8). The marginal likelihood of each model is computed as

$$p(\mathbf{R}|\mathbf{F};\mathcal{M}_i) = \int ... \int \sum_{\mathcal{K}} p(\mathbf{R}|\mathbf{B},\mathcal{K},\mathbf{\Sigma},\boldsymbol{\lambda},\boldsymbol{\theta},\mathbf{F};\mathcal{M}_i) \times p(\boldsymbol{\theta},\mathbf{B},\mathcal{K},\mathbf{\Sigma},\boldsymbol{\lambda}|\mathbf{R},\mathbf{F};\mathcal{M}_i)d\mathbf{B}d\mathbf{\Sigma}d\boldsymbol{\theta}d\mathbf{\Sigma}, \tag{16}$$

where $\mathcal{M}_i$ identifies the ith model and the posterior density $p(\mathbf{R},\mathbf{B},\mathcal{K},\mathbf{\Sigma},\boldsymbol{\lambda}|\boldsymbol{\theta},\mathbf{F};\mathcal{M}_i)$ is given by (15). Following Chib (1995), we compute the marginal likelihood by replacing the unobservable breaks and parameters in the likelihood of the data generating process defined by (5)- (8) for each draw.

Table 3 reports the marginal (log)likelihoods for each of the model specifications as well as the Bayes factors, the difference between model-specific (log)likelihoods, used as an model selection indicator that naturally penalizes for the different size/complexity of different models (see Kass and Raftery, 1995), for each of the alternative frameworks including the two-step F-MB approach, vs. B-TVB-SV. Because Bayes factors are constructed from marginal likelihoods, they measure a model ability to explain the entire distribution (not just first moments) of test asset returns. Bayes factors also permit the simultaneous comparison of multiple models, regardless of whether the models are nested. To favor interpretations, also the (log)likelihood contributions by each of the 23 test portfolios under each of the models and the corresponding Bayes factors have been computed. Interestingly, the B-TVB-SV model shows the higher marginal (log) likelihood values across all of the portfolios under consideration, as well as the higher overall marginal likelihood. By exceeding 100, all the overall Bayes factors are highly significant. In particular, the factors vs. the B-TVP and the two-step F-MB implementations are 892 and 5602, respectively, and therefore appear to be decisively in favor of the complete B-TVB-SV framework. The Bayes factor vs. the B-TVB model with stochastic volatility is instead 191 and remains favorable to B-TVB-SV. Surprisingly, the B-TVP model ranks second both in overall terms and for all the test portfolios, thus outperforming the homoskedastic B-TVB alternative. This result emphasizes that by fully acknowledging instability in idiosyncratic risk plays a key role beyond that of capturing breaks in the betas. As one would expect, given its ingenious but ad-hoc nature, the classical two-step F-MB approach ranks last with an overall marginal likelihood around 15 times lower than under the B-TVB-SV model. The dominance of the B-TVB-SV framework occurs across all portfolios, but appears to be particularly elevated in the case of bonds and medium and large caps portfolios of stocks.

4.6. *Robustness Checks*

We have experimented with an informative prior in the second pass in order to put some structure (constraints) on the distribution and moments of the risk premia. These are now postulated to be normally distributed with zero mean and variance such that there is 95% probability that annualized premia are smaller in absolute value than the largest between the absolute value between the minimum and the maximum excess return observed in our sample.[19] We record a considerable reduction in the variability of the estimated posterior distributions of the risk premia relative to the baseline case. The qualitative results and insights from Table 2 in the B-TVB-SV case apply intact and, in general, using informative priors on the premia limits their variability so we find both more precisely estimated premia (so far the result has been built in the type of prior used) and economic implications that encompass Tables 2 and 3: industrial production growth, liquidity, unexpected inflation, and especially market risks

[19] A complete description of prior distributions and hyperparameters used can be found in Appendix A.

are important drivers of the cross section of U.S. stock and bond returns. To save space, we have not plotted or tabulated complete set of results, that remain available from the Authors.

5. **Economic Assessment**

So far our discussion has focussed on statistical performance in terms of whether there was evidence of either the $\lambda_{0,t}$s or the $\beta_{i0,t}$s coefficients being different from zero and especially with emphasis on the comparison of marginal log-likelihood values. We have concluded that (1)-(2) is rejected in its two-pass F-MB implementation based on 5-year rolling window estimates. However, there was some supportive indications that the B-TVB-SV model may be not completely at odds with the data. The results concerning B-TVP have shown that while there are some degrees of freedom as to the way one ought to best model time-variation in risk exposures, capturing instability in stochastic volatility is truly fundamental. Yet, we still know little about the economic implications of B-TVB-SV. In this section, we report additional evidence on the economic importance of the estimates uncovered for B-TVB-SV model. In Section 5.1, we comment on the variance ratios, $VR1$ and $VR2$ described in Appendix B that measure the degree of misspecification of a MFAPM. The idea of $VR1$ and $VR2$ is that a correctly specified MFAPM should at least explain most or all of the predictable variation in the excess returns of the test assets, and therefore leave an unexplained portion that should be as small as possible. Using results detailed in Appendix B, Section 5.2 comments on a decomposition of the sources of predictable variation of excess returns due to the MFAPM. Section 5.3 reports pricing tests that provide an alternative and intuitive measure of the quality of the approximation provided by the MFAPM.

5.1. *Variance Ratios*

With reference to the estimates of (5)-(8), we have computed (posterior distributions of the) VR1 and VR2 ratios defined in Appendix B and typical of the literature. Given their popularity, we just limit ourselves to recall that 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 macroeconomic risk; at the same time, VR2 should be equal to zero if the multi-factor model is correctly specified. Moreover, as explained in Appendix B, $VR1 = 1$ does not imply that $VR2 = 0$ and viceversa. In what follows, the information at time $t-1$ ($\mathbf{Z}_{t-1}$) used to tease out the total predictable variation in excess returns used as a "denominator" in the empirical results that follow is proxied by the instrumental variables listed in Table 1, plus a dummy variable to account for the so-called "January effect" (see Thaler, 1987).

Columns 4 and 7 of Table 4 present posterior medians of (normalized) $VR1$ and $VR2$ obtained from the B-TVB-SV model for each of the 23 portfolios. These variance ratios are compared to the ones obtained from competing models. A normalization is performed by dividing the posterior medians by the variance of the underlying excess return series. Variance ratio results are encouraging. Under a VR1 perspective, on average approximately 80% of the predictable variation in excess returns is captured by the B-TVB-SV model. Such a statistic is only 51% in the case of the F-MB implementation (column 1) and goes as low as 47 and 43% for the B-TVP and homoskedastic B-TVB models, respectively. Although in the light of the earlier marginal log-likelihood evidence, this is relatively un-surprising, because the mapping between the ability to capture any predictable variation and the log-likelihood is a complex one, these results remain economically meaningful. However, the generally high VR1 ratios from the B-TVB-SV model vary considerably across different test assets. The ratios are relatively high, also in relation to what is typically reported in the literature (see Ferson and Harvey, 1991, or more recently Guidolin et al., 2013), in the case of government bond portfolios (possibly because we have used Cochrane and Piazzesi's factor) and for a few industries, such as manufacturing, energy, and high-tech, for which VR1 exceeds 90%. It is instead below 50% in the case of the smallest capitalization decile and of non-investment grade corporate bonds, exactly where one would expect our macroeconomic risk factors to have more trouble at fitting the variation in excess returns.[20]

Because $VR1 + VR2 = 1$ does not hold, the finding of high VR1 ratios fails to imply that the VR2 ratios are close to zero. Yet, VR2 is on average just above 20% in the case B-TVB-SV, to be contrasted with averages across test portfolios of 48-54% in the case of other models. Moreover, in the case of the B-TVB-SV framework, we record VR2 ratios equal to or inferior to 15% in 9 out of 23 portfolios. All in all, under both the VR1 and VR2, we find evidence of appreciable performance of the model.

5.2. *Sources of Risk*

We have also followed Appendix B and computed the contribution of each factor to the fit offered by the B-TVB-SV model to fitting the predictable variation in excess stock returns. The highest contribution is given by the market risk factor: with three exceptions (energy, health, and utility stocks), all the ratios $Var[P(\lambda_{MKT,t}\beta_{iMKT,t|t-1}|\ \mathbf{Z}_{t-1})]\ /Var[P(\sum_{j=1}^{9}\lambda_{j,t}\beta_{ij,t|t-1}|\mathbf{Z}_{t-1})]$ concerning stocks exceed 0.5 with peaks in excess of 1 for a number of industries as well as medium-cap portfolios.[21] However, the market factor does not explain most of the predictable variation in excess bond returns, when it is replaced in this leading role by the credit risk factor. As far as stocks are concerned, the next most important contributions come from unexpected inflation (especially for bond and selected industry portfolios) and to some extent, real consumption growth risk, although the heterogeneity across portfolios is pronounced (small capitalization stocks are particularly well explained by this factor). In the case of bond portfolios, most predictable variation is explained, after taking into credit risk exposures into account, by unexpected inflation, economy-wide market risks, and Cochrane and Piazzesi's specific factor;

[20] The 95% credibility regions do little to cast any doubts on this obvious outperformance of the B-TVB-SV framework over the competing models. For instance, the 2.5% posterior lower bound in the case of B-TVB-SV is 62.2% vs. 22.6% in the case of B-TVP and 19.0% in the case of the homoskedastic B-TVB.

[21] These ratios may exceed 100% because $Var[P(\sum_{j=1}^{9}\lambda_{j,t}\beta_{ij,t|t-1}|\mathbf{Z}_{t-1})]$ will also reflect the contribution of covariance terms between factor terms.

interestingly, the contribution of yield curve shocks to priced risk is limited.

5.3. *Pricing Errors*

We follow Geweke and Zhou (1996) and measure the closeness of the pricing approximation provided by an approximate version of (9), $E_{t-1}[r_{i,t}] \simeq \lambda_{0,t} + \sum_{j=1}^{K} \beta_{ij,t}\lambda_{j,t}$, by computing at each time t the average squared recursive pricing error across all the N test assets/portfolios,

$$Q^2_{t,N} = \frac{1}{N}\left[\boldsymbol{\beta}'_{0,t}\left(\mathbf{I}_N - \mathbf{B}_t\left(\mathbf{B}'_t\mathbf{B}_t\right)^{-1}\mathbf{B}'_t\right)\boldsymbol{\beta}_{0,t}\right] \qquad t = 1, ..., T, \tag{17}$$

where $\boldsymbol{\beta}_{0,t}$ is the $N \times 1$ vector of intercepts, $\mathbf{I}_N$ is an N-dimensional identity matrix, and $\mathbf{B}_t \equiv (\boldsymbol{\iota}_N, \boldsymbol{\beta}_{1,t}, ..., \boldsymbol{\beta}_{K,t})$ is a $N \times K$ matrix collecting vectors of time t betas of all the assets/portfolios vs. each of the K risk factors, with $\boldsymbol{\beta}_{j,t} \equiv (\beta_{1j,t}, ..., \beta'_{Nj,t})'$ a $N \times 1$ vector of factor loadings on the jth risk factor. These pricing errors are recursive because at each point in time they are obtained using only information available up to that point. Because our Gibbs sampling scheme allows to derive posteriors for all the objects that enter $\boldsymbol{\beta}_{0,t}$ and $\mathbf{B}_t$, we also compute the posterior density of the average (squared) pricing error statistic, as discussed in Geweke and Zhou (1996).

Table 5 reports the average monthly pricing errors, $Q_{t,N}$, for each of our models across different sub-samples. Using sub-samples wants to allow any instability in pricing performance to emerge and be adequately detected.[22] With reference to the full-sample, the B-TVB-SV model yields both the lowest average pricing error (0.21% per month) and the lowest median posterior error (0.19%). Such statistics are practically between one-half and two-thirds those that one would obtain under a B-TVB homoskedastic model (0.41 and 0.35 percent, respectively). Interestingly, the B-TVP model seems to fit the data well on the basis of the Bayes odds ratio in Table 3, but fails to price our test portfolios (it gives average and median posterior errors of 0.54 and 0.51 percent, respectively) as accurately as the homoskedastic B-TVB model does. The performance of the classical two-step F-MB scheme is poor, yielding average and median pricing errors of 0.63%. Moreover, B-TVB-SV consistently outperforms all other models in all sub-samples. Its advantage is always substantial in the sense that B-TVB-SV always cuts the average error of the second best model by at least 40%. Interestingly, the pricing errors tend to increase over our sample, especially when one compares the 1982-1998 with the 1999-2011 interval. However, there is no evidence of the errors during the Great Financial Crisis period (2007-2011) being systematically higher than in the overall 1999-2011 sub-sample.

Figure 11 plots the time series of average pricing errors $Q_{t,N}$ for all the models (top panel) and only for the B-TVB-SV and the B-TVP cases, rescaling the errors from the former model (on the right axis) to better emphasize similarities and differences (bottom panel). The top panel shows that, apart

[22] The 2007-2011 sub-sample specifically addresses pricing performance during the recent financial crisis.

from a short period in early 1993 (when B-TVB became competitive), B-TVB-SV gave uniformly lower average pricing errors than all other models. The F-MB scheme gives uniformly high but constant average errors. The B-TVP model gives a highly variable performance, with enormous spikes of misspricing around 1993, in 1999-2000, and during the financial crisis. The bottom panel of the figure shows that the dynamics of pricing errors under B-TVB-SV and B-TVP—both models including a stochastic volatility component—are not that dissimilar, in the sense that also errors from B-TVB-SV spike up in 1993, 1999-2001 and during 2011. However, the more parsimonious dynamics imposed by infrequent, large structural breaks under B-TVB-SV reduces the pricing errors also keeping the latter more stable across our sample.

6. Conclusion

In this paper, we have proposed a new way to parameterize and estimate in state-space form a typical MFAPM with time-varying risk exposures and premia. This Bayesian state-space approach is based on a formal modelling of the latent process followed by risk exposures and idiosyncratic volatility capable to capture structural shifts in parameters. This method can also be interpreted as a novel way to overcome the two-pass approach advocated by Fama and MacBeth (1973) and used in a substantive body of applied work in finance. Given a general B-TVB-SV framework, we have also considered special cases that are obtained by imposing restrictions, and in particular a B-TVP-SV model in which betas change continuously but in small amounts, and a homoskedastic B-TVB model in which volatility is constant.

Our application to monthly, 1972-2011 U.S. stock and bond returns shows that the two-stage approach yields results that are not always reasonable. For instance, very few risk factors appear to be priced and, when they are, they carry the wrong sign. Moreover, the fit provided by the standard two-step approach is poor. On the contrary, the empirical implications of a Bayesian state-space implementation of (5)- (6) are plausible and there are indications that the model is consistent with the data. For instance, most portfolios do not appear to have been grossly mispriced and a few risk premia are precisely estimated with a plausible sign. Market, liquidity, and industrial production (real output) growth risks are significantly priced. This confirms the early evidence in McElroy and Burmeister (1988) that appropriate econometric methods reveal a strong explanatory power of macroeconomic factors in addition to that provided by the plain vanilla, CAPM-style market portfolio. Bayes odds ratios and marginal likelihood comparisons indicate that the B-TVB-SV outperforms both the two-step F-MB and the homoskedastic B-TVB models. The heteroskedastic B-TVP appears to be closer to the full-scale B-TVB-SV one. However, an analysis of the average pricing errors shows that large but infrequent breaks in factor exposures are considerably more successful. Finally, the finding that the heteroskedastic B-TVB models ranks second below B-TVB-SV is a powerful indication of the importance to explicitly model stochastic volatility when implementing multi-factor asset pricing models.

References

[1] Aguilar O., and West, M., "Bayesian Dynamic Factor Models and Portfolio Allocation", *Journal of Business and Economic Statistics*, 2000, 18, 338-357.

[2] Ang A. and Chen J., "CAPM Over the Long Run: 1926-2001", *Journal of Empirical Finance*, 2007, 14, 1-40.

[3] Bekaert G., Hodrick R.J., and Zhang X., "Aggregate Idiosyncratic Volatility", *Journal of Financial and Quantitative Analysis*, 2012, 47, 1155-1185.

[4] Bossaerts P., and Green, R., "A General Equilibrium Model of Changing Risk Premia: Theory and Tests", *Review of Financial Studies*, 1989, 2, 467-493.

[5] McElroy M.B. and Burmeister E., "Arbitrage Pricing Theory as a Restricted Nonlinear Multivariate Regression Model", *Journal of Business Economics and Statistics*, 1988, 6, 29-42.

[6] Campbell J.Y., Lettau M. , Malkiel B.G. and Xu Y., "Have individual stocks become more volatile? An empirical exploration of idiosyncratic risk", *Journal of Finance*, 2001, 56, 1-43.

[7] Campbell J.Y. and Cochrane J.H., "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior", *Journal of Political Economy*, 1999, 107, 205-251.

[8] Carter C., and Kohn R., "On the Gibbs Sampling for State-Space Models", *Biometrika*, 1994, 81, 541-553.

[9] Carvalho, C.M., Lopes, H.F., and Aguilar, O., "Dynamic Stock Selection Strategies: A Structured Factor Model Framework", *Bayesian Statistics*, 9, 2011, 1-21.

[10] Chan L.K.C., Karceski J., and Lakonishok J., "The Risk and Return from Factors", *Journal of Financial and Quantitative Analysis*, 1998, 33, 159-188.

[11] Chen N.-F., Roll R. and Ross S., "Economic Forces and the Stock Market", *Journal of Business*, 1986, 59, 383-403.

[12] Chib, S., "Marginal Likelihood from the Gibbs Output", *Journal of the American Statistical Association*, 1995, 432, 1313-1321.

[13] Cochrane J.H., *Asset Pricing*, Princeton University Press, Princenton: NJ, 2005.

[14] Cochrane J.H., and Piazzesi, M. "Bond Risk Premia", *American Economic Review*, 2005, 94, 138-160.

[15] De Pooter M., Ravazzolo F., Segers R. and van Dijk H.K., "Bayesian Near-Boundary Analysis in Basic Macroeconomic Time Series Models", *Advances in Econometrics*, 2008, 23, 331-402.

[16] Dittmar R.F., "Nonlinear Pricing Kernels, Kurtosis Preference, and Evidence from the Cross Section of Equity Returns", *Journal of Finance*, 2002, 57, 369-403.

[17] Fama E., and MacBeth J., "Risk, Return and Equilibrium: Empirical Tests", *Journal of Political Economy*, 1973, 81, 607-636.

[18] Ferson W., and Harvey C. R., “The Variation of Economic Risk Premiums”, *Journal of Political Economy*, 1991, 99, 385-415.

[19] Flannery M.J., and Protopapadakis A.A., “Macroeconomic Factors Do Influence Aggregate Stock Returns”, *Review of Financial Studies*, 2002, 15, 751-782.

[20] Frühwirth-Schnatter, S., “Data Augmentation and Dynamic Linear Models”, *Journal of Time Series Analysis*, 1994, 15, 183-202.

[21] Geman S., and Geman D., “Stochastic Relaxation, Gibbs Distributions, and the Bayesian Restoration of Images” *IEEE Transactions*, 1984, 6, 721-741.

[22] George E., McCulloch R.E., “Variable Selection via Gibbs Sampling”, *Journal of the American Statistical Association*, 1993, 88, 881-889.

[23] Gerlach R., Carter C. and Kohn R., “Efficient Bayesian Inference for Dynamic Mixture Models”, *Journal of the American Statistical Association*, 2000, 95, 819-828.

[24] Geweke J. and Zhou, G. “Measuring the Pricing Error of the Arbitrage Pricing Theory”, *Review of Financial Studies*, 1996, 9, 557-587.

[25] Giordani P., Kohn R. and van Dijk D., “A Unified Approach to Nonlinearity, Outliers and Structural Breaks”, *Journal of Econometrics*, 2007, 137, 112-133.

[26] Giordani P. and Kohn R., “Efficient Bayesian Inference for Multiple Change-Point and Mixture Innovation Models”, *Journal of Business and Economic Statistics*, 2008, 26, 1, 66-77.

[27] Giordani P., and Villani, M., “Forecasting Macroeconomic Time Series with Locally Adaptive Signal Extraction”, *International Journal of Forecasting*, 2010, 26, 312-325.

[28] Groen J.J.J., Paap R. and Ravazzolo, F., “Real-time Inflation Forecasting in a Changing World”, *Journal of Business and Economic Statistics*, 2013, 31, 29-44.

[29] Guidolin M., Ravazzolo F., and Tortora A.D., “Alternative Econometric Implementations of Multi-Factor Models of the U.S. Financial Markets”, *Quarterly Review of Economics and Finance*, 2013, 53, 87-111.

[30] Gungor S., and Luger, R., “Testing Linear Factor Pricing Models With Large Cross Sections: A Distribution-Free Approach”, *Journal of Business and Economic Statistics*, 2013, 31, 66-77.

[31] Harvey C.R., “The Real Term Structure and Consumption Growth, *Journal of Financial Economics*, 1988, 22, 305-333.

[32] Ludvigson, S.C., and Ng S., “Macro Factors in Bond Risk Premia”, *Review of Financial Studies*, 2009, 22, 5027-5067.

[33] Jostova G. and Philipov A., “Bayesian Analysis of Stochastic Betas”, *Journal of Financial and Quantitative Analysis*, 2005, 40, 747-778.

[34] Karolyi G.A. and Sanders A., “The Variation of Economic Risk Premiums in Real Estate Returns”, *Journal of Real Estate Finance and Economics*, 1998, 17, 245-262.

[35] Kass R.E., and Raftery A., "Bayes Factors", *Journal of the American Statistical Association*, 1995, 430, 773-795.

[36] Kim S., Shepard N. and Chib S., "Stochastic volatility: likelihood inference and comparison with ARCH models", *Review of Economic Studies*, 1998, 65, 361-93.

[37] Koop, G., and Potter, S.M., "Estimation and Forecasting in Models with Multiple Breaks", *Review of Economic Studies*, 2007, 74, 763-789.

[38] Kramer, C., "Macroeconomic Seasonality and the January Effect", *Journal of Finance*, 1994, 49, 1883-1891.

[39] Lamont O., "Economic Tracking Portfolios", *Journal of Econometrics*, 2001, 105 161-184.

[40] Lopes, H.F., and Carvalho, C.M., "Factor Stochastic Volatility with Time-Varying Loadings and Markov Switching Regimes", *Journal of Statistical Planning and Inference*, 2007, 137, 3082-3091.

[41] Maheu, J.M., and Gordon, S., "Learning, Forecasting and Structural Breaks", *Journal of Applied Econometrics*, 2008, 23, 553-583.

[42] Maheu, J.M., and McCurdy, T.H., "How Useful are Historical Data for Forecasting the Long-Run Equity Return Distribution?", *Journal of Business and Economic Statistics*, 2009, 27, 95-112.

[43] McCulloch R., and Rossi P., "A Bayesian Approach to Testing the Arbitrage Pricing Theory", *Journal of Econometrics*, 1991, 49, 141-168.

[44] McQueen G., and Roley V.V., "Stock Prices, News, and Business Conditions", *Review of Financial Studies*, 1993, 6, 683-707.

[45] Merton R., "An Intertemporal Capital Asset Pricing Model", *Econometrica*, 1973, 41, 867-887.

[46] Miazhynskaia, T., Frühwirth-Schnatter, S, and Dorffner, G., "Bayesian testing for non-linearity in volatility modeling", *Computational Statistics & Data Analysis*, 2006, 51, 2029-2042.

[47] Mitchell, T., and Beauchamp, J., "Bayesian Variable Selection in Linear Regression", *Journal of the American Statistical Association*, 1988, 83, 1023-1032.

[48] Næs R., Skjeltorp J.A., and Ødegaard B.A., "Stock Market Liquidity and the Business Cycle", *Journal of Finance*, 2011, 66, 139-176.

[49] Omori, Y., Chib S., Shepard N., and Nakajima, J., "Stochastic Volatility with Leverage: Fast and Efficient Likelihood Inference", *Journal of Econometrics*, 2007, 140, 425-449.

[50] Pagan, A., "Econometric Issues in the Analysis of Regressions with Generated Regressors", *International Economic Review*, 25, 221-247.

[51] Pastor, L. and Stambaugh, R.F., "Liquidity Risk and Expected Stock Returns", *Journal of Political Economy*, 2003, 111, 642-685.

[52] Petersen, M.A., "Estimating Standard Errors in Finance Panel Data Sets: Comparing Approaches", *Review of Financial Studies*, 2008, 22, 435-480.

[53] Primiceri G., "Time Varying Structural Vector Autoregressions and Monetary Policy", *Review of Economic Studies*, 2005, 72, 821-852.

[54] Ravazzolo F., Paap R., van Dijk D. and Franses P.H., "Bayesian Model Averaging in the Presence of Structural Breaks", Rapach D. and Wohar M. (eds.), "Forecasting in the Presence of Structural Breaks and Uncertainty", *Frontiers of Economics and Globalization*, Elsevier, 2007.

[55] Ross S., "The Arbitrage Theory of Capital Asset Pricing" *Journal of Economic Theory*, 1976, 13, 341-360.

[56] Shanken J., "On the Estimation of Beta Pricing Models", *Review of Financial Studies*, 1992, 5, 1-34.

[57] Shanken J., and Weinstein M., "Economic Forces and the Stock Market Revisited", *Journal of Empirical Finance*, 2006, 13, 129-144.

[58] Shiller R.J., "The Volatility of Long-Term Interest Rates and Expectations Models of the Term Structure", *Journal of Political Economy*, 1979, 87, 1190-1219.

[59] Singleton, K., *Empirical Dynamic Asset Pricing: Model Specification and Econometric Assessment*, Princeton University Press, Princenton: NJ, 2006.

[60] Tanner M.A., and Wong W.H., "The Calculation of Posterior Distributions by Data Augmentation", *Journal of the American Statistical Association*, 1987, 82, 528-550.

[61] Thaler R.H., "Anomalies: The January Effect", *Journal of Economic Perspectives*, 1987, 1, 197-201.

[62] West, M., and P.J., Harrison, *Bayesian Forecasting and Dynamic Models*, New York: Springer-Verlag, 2007.

Appendix A: The Gibbs Sampling Algorithm

Here we derive the full conditional posterior distributions of the latent variables and the model parameters discussed in Section 3. For the ease of exposition we report the results for the ith asset.

Step 1: Sampling K_β.

The structural breaks in the conditional dynamics of the factor loadings measured by the latent binary state $\kappa_{i0t,...,}\kappa_{iKt}$, are drawn using the algorithm of Gerlach et al. (2000). This algorithm increases the efficiency of the sampling procedure since allows to generate $\boldsymbol{\kappa}_{it} = (\kappa_{i0t,...,}\kappa_{iKt})$, without conditioning on the relative regression parameters $\boldsymbol{\beta}_{it} = (\beta_{i0t}, ..., \beta_{iKt})$. The conditional posterior density of $\boldsymbol{\kappa}_{it}$, $t = 1, ..., T, i = 1, ..., N$, for each of ith asset/portfolio is defined as

$$\begin{aligned} p\left(\kappa_{i0t,...,}\kappa_{iKt}|\mathcal{K}_{i\beta[-t]}, \mathcal{K}_{i\sigma}, \mathbf{\Sigma}_i, \theta_i, \mathbf{R}_i, \mathbf{F}\right) &\propto p\left(\mathbf{R}_i|\mathcal{K}_{i\beta}, \mathcal{K}_{i\sigma}, \mathbf{\Sigma}_i, \theta_i, \mathbf{F}\right) p\left(\kappa_{i0t,...,}\kappa_{iKt}|\mathcal{K}_{i\beta[-t]}, \mathcal{K}_{i\sigma}, \mathbf{\Sigma}_i, \theta_i, \mathbf{F}\right) \\ &\propto p\left(r_{i,t+1,...,}r_{i,T}|r_{i,1,...,}r_{i,t}, \mathcal{K}_{i\beta}, \mathcal{K}_{i\sigma}, \mathbf{\Sigma}_i, \theta_i, \mathbf{F}\right) p\left(r_{i,t}|r_{i,1,...,}r_{i,t-1}, \mathcal{K}_{i\beta[1:t]}, \mathcal{K}_{i\sigma}, \mathbf{\Sigma}_i, \theta_i, \mathbf{F}\right) \\ &\qquad p\left(\kappa_{i0t,...,}\kappa_{iKt}|\mathcal{K}_{i\beta[-t]}, \mathcal{K}_{i\sigma}, \mathbf{\Sigma}_i, \theta_i, \mathbf{F}\right) \end{aligned} \tag{18}$$

where $\mathcal{K}_{i\beta[-t]} = \left\{ \{\kappa_{ijs}\}_{j=0}^{K} \right\}_{s=1,s\neq t}^{T}$, $\mathcal{K}_{i\beta[1:t]} = \left\{ \{\kappa_{ijd}\}_{j=0}^{K} \right\}_{d=1}^{t}$ and $\mathcal{K}_{i\sigma} = \{\kappa_{i\nu,t}\}_{t=1}^{T}$. We assume that each of the κ_{ijs} breaks are independent from each other such that the joint density is defined as $\prod_{j=0}^{K} \pi_{ij}^{\kappa_{ijs}} (1-\pi_{ij})^{1-\kappa_{ijs}}$. The remaining densities $p(r_{i,t+1,...,}r_{i,T}|r_{i,1,...,}r_{i,t}, \mathcal{K}_{i\beta}, \mathcal{K}_{i\sigma}, \mathbf{\Sigma}_i, \theta_i, \mathbf{F})$ and $p(r_{it}|r_{i1,...,}r_{it-1}, \mathcal{K}_{i\beta t}, \mathcal{K}_{i\sigma}, \mathbf{\Sigma}_i, \theta_i, \mathbf{F})$ are evaluated as in Gerlach et al. (2000). Notice that, since κ_{ijt} is a binary state the integrating constant is easily evaluated.

Step 2: Sampling the Factor Loadings $\mathbf{B}$.

The full conditional posterior density for the time-varying factor loadings is computed using a standard forward filtering backward sampling as in Carter and Kohn (1994). For each of the $i = 1, ..., N$ assets, the prior distribution of the $\beta_{i0}, ..., \beta_{iK}$ loadings is a multivariate normal with the location parameters corresponding to the OLS parameter estimates and a covariance structure which is diagonal and defined by the variances of the OLS estimates. The initial prior are sequentially updated via the Kalman Filtering recursion, then the parameters are drawn from the posterior distribution which is generated by a standard backward recursion (see Fruhwirth-Schnatter 1994, Carter and Kohn 1994, and West and Harrison 1997).

Steps 3 and 4: Sampling the Breaks and the Values of the Idiosyncratic Volatility.

In order to draw the structural breaks $\mathcal{K}_{i\sigma}$ and the idiosyncratic volatilities $\mathbf{\Sigma}_i$ for each of the ith portfolios, we follow a similar approach as in step 1. The stochastic breaks $\mathcal{K}_{i\sigma}$ are drawn by using the Gerlach et al. (2000) algorithm. The conditional variances $\ln \sigma_{it}^2$, does not show a linear structure even though still preserving the standard properties of state space models. The model is rewritten as

$$\ln \left(r_{i,t} - \beta_{i0t} - \sum_{j=1}^{K} \beta_{ijt} F_{jt} \right)^2 = \ln \sigma_{it}^2 + u_t$$

$$\ln \sigma_{it}^2 = \ln \sigma_{it-1}^2 + \kappa_{\nu it} \nu_{it} \tag{19}$$

where $u_t = \ln \varepsilon_t^2$ has a $\ln \chi^2(1)$. Here we follow Omori et al. (2010) and 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

$$p(u_t) = \sum_{l=1}^{10} \varphi_l \frac{1}{\sqrt{\varpi_l^2 2\pi}} \exp\left(-\frac{(u_t - \mu_l)^2}{2\varpi_l} \right) \tag{20}$$

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. (2010). Mechanically, in each step of the sampler we simulate at each time t a component of the mixture. Given the mixture component, we can apply the standard Kalman filter such that $\mathcal{K}_{i\sigma}$ and $\mathbf{\Sigma}_i$ can be sampled in a similar way as $\mathcal{K}_{i\beta[t]}$ and $\beta_{i0[t]}, ..., \beta_{iK[t]}$ in the first and second step. The initial prior of the log idiosyncratic volatility $\ln \sigma_{i0}^2$ is normal with mean -1 and conditional variance equal to 0.1.

Step 5b. Sampling the Time-Varying Risk Premia .

The cross-sectional equilibrium restriction in (2) is satisfied at each time t conditional on the latent states $\mathbf{B}_{t|t-1} = \left\{ \left\{ \beta_{ijt|t-1} \right\}_{i=1}^{N} \right\}_{j=0}^{K}$ and $\mathbf{\Sigma}_t = \left\{ \boldsymbol{\sigma}_{it}^2 \right\}_{i=1}^{N}$. Given an initial normal-inverse gamma prior, the full conditional of the equilibrium risk premia $\boldsymbol{\lambda}_t = (\lambda_{0t}, ..., \lambda_{Kt})$ at time t, is defined as

$$p(\boldsymbol{\lambda}_t|\tau, \mathbf{B}_{t|t-1}, \mathbf{\Sigma}_t, \mathbf{R}_t) \propto |\mathbf{\Sigma}_t^*|^{-\frac{1}{2}} \exp\left\{ -\frac{1}{2} \left(\mathbf{R}_t - \boldsymbol{\mu}_t^*\right)' \left(\mathbf{\Sigma}_t^*\right)^{-1} \left(\mathbf{R}_t - \boldsymbol{\mu}_t^*\right) \right\} \tag{21}$$

where $\mathbf{R}_t = (r_{1t,...,}r_{N,t})$ and $\mathbf{\Sigma}_0$, $\boldsymbol{\mu}_0$ respectively the prior mean and variance of λ_t, such that the conditional (ex-ante time-varying) risk premia can be sampled at each time t by a normal distribution with $\boldsymbol{\mu}_t^* = \mathbf{\Sigma}_t^*(\mathbf{\Sigma}_0^{-1}\boldsymbol{\mu}_0 + \tau^{-2}\mathbf{X}_{t-1}'\mathbf{R}_t)$ and $\mathbf{\Sigma}_t^* = (\mathbf{\Sigma}_0^{-1} + \tau^{-2}\mathbf{X}_{t-1}'\mathbf{X}_{t-1})^{-1}$, $\mathbf{X}_{t-1} \equiv \left[\iota, \mathbf{B}_{t|t-1}\right]$, respectively as location and scale parameters. The conditional posterior for the variance of the risk premia τ^2 is an inverse gamma distribution

$$p(\tau^2|\boldsymbol{\lambda}_t, \mathbf{B}_{t|t-1}, \mathbf{\Sigma}_t, \mathbf{R}_t) \propto \tau^{-a_0} \exp\left(-\frac{b_0}{2\tau}\right) \prod_{i=1}^{N} \frac{1}{\tau} \exp\left(\frac{-\left(r_{it} - \lambda_{0t} - \sum_{j=1}^{K} \beta_{ijt|t-1}\lambda_{jt}\right)^2}{2\tau^2} \right) \tag{22}$$

such that τ^2 can be sampled from an inverse-gamma distribution with scale parameter $b = b_0 + \sum_{i=1}^{N} \left(r_{it} - \lambda_{0t} - \sum_{j=1}^{K} \beta_{ijt|t-1}\lambda_{jt}\right)^2$ and degrees of freedom $a = a_0 + N$.

Step 5b. Sampling the Stochastic Breaks Probabilities.

The full conditional posterior densities for the breaks probabilities $\boldsymbol{\pi} = (\pi_{i1,...,}\pi_{iK})$ is given by

$$p\left(\boldsymbol{\pi}|q^2, \mathbf{B}, \mathbf{\Sigma}, \mathcal{K}_\beta, \mathbf{R}, \mathbf{F}\right) \propto \prod_{j=0}^{K} \pi_{ij}^{a_{ij}-1} \left(1 - \pi_{ij}\right)^{b_{ij}-1} \prod_{t=1}^{T} \pi_{ij}^{\kappa_{ijt}} (1 - \pi_{ij})^{1-\kappa_{ijt}} \tag{23}$$

and hence the individual π_{ij} parameter can be sampled from a Beta distribution with shape parameters $a_{ij} + \sum_{t=1}^{T} \kappa_{ijt}$ and $b_{ij} + \sum_{t=1}^{T}(1 - \kappa_{ijt})$ for $j = 0, ..., K$. Likewise the full conditional posterior distribution for the breaks probabilities in the idiosyncratic volatilities π_ν is given by

$$p\left(\pi_\nu|q^2, \mathbf{B}, \mathbf{\Sigma}, \mathcal{K}_\sigma, \mathbf{R}, \mathbf{F}\right) \propto \pi_{i\nu}^{a_{i\nu}-1} \left(1 - \pi_{i\nu}\right)^{b_{i\nu}-1} \prod_{t=1}^{T} \pi_{i\nu}^{\kappa_{i\nu t}} (1 - \pi_{i\nu})^{1-\kappa_{i\nu t}}$$

such that the individual $\pi_{i\nu}$ can be sampled from a Beta distribution with shape parameters $a_{i\nu} + \sum_{t=1}^{T} \kappa_{i\nu t}$ and $b_{i\nu} + \sum_{t=1}^{T}(1 - \kappa_{i\nu t})$ for $i = 1, ..., N$.

Step 5c. Sampling the Conditional Variance of the States.

The prior distributions for the conditional volatilities of the factor loadings β_{ijt} for $j = 0, ..., K$ are inverse-gamma

$$p\left(q_{ij}^2|\pi, \mathbf{B}, \mathbf{\Sigma}, \mathcal{K}_\beta, \mathcal{K}_\sigma, \mathbf{R}, \mathbf{F}\right) \propto q_{ij}^{-\nu_{ij}} \exp\left(-\frac{\delta_{ij}}{2q_{ij}^2}\right) \prod_{t=1}^{T} \left(\frac{1}{q_{ij}} \exp\left(\frac{-\left(\beta_{ijt} - \beta_{ijt-1}\right)^2}{2q_{ij}^2} \right) \right)^{\kappa_{ijt}} \tag{24}$$

hence q_{ij}^2 is sampled from an inverse-gamma distribution with scale parameter $\nu_{ij} + \sum_{t=1}^{T} \kappa_{ijt}(\beta_{ijt} -$

$\beta_{ijt-1})^2$ and degrees of freedom equal to $\nu_{ij} + \sum_{t=1}^{T} \kappa_{ijt}$. Likewise the full conditional of the variance for the idiosyncratic log volatility $q_{i\nu}^2$ is defined as

$$p\left(q_{i\nu}^2 | \pi, \mathbf{B}, \mathbf{\Sigma}, \mathcal{K}_\beta, \mathcal{K}_\sigma, \mathbf{R}, \mathbf{F}\right) \propto q_{i\nu}^{-\nu_{i\nu}} \exp\left(-\frac{\delta_{i\nu}}{2q_{i\nu}^2}\right) \prod_{t=1}^{T} \left(\frac{1}{q_{i\nu}} \exp\left(\frac{-\left(\ln\sigma_{it}^2 - \ln\sigma_{it-1}^2\right)^2}{2q_{i\nu}^2}\right)\right)^{\kappa_{i\nu t}} \quad (25)$$

such that $q_{i\nu}^2$ is sampled from an inverted Gamma distribution with scale parameter $\nu_{i\nu} + \sum_{t=1}^{T} \kappa_{i\nu t}(\ln\sigma_{it}^2 - \ln\sigma_{it-1}^2)^2$ and degrees of freedom equal to $\nu_{i\nu} + \sum_{t=1}^{T} \kappa_{i\nu t}$.

Choice of Priors

Realistic values for the different prior distributions obviously depend on the problem at hand.[23] In general, we use weak priors, excluding the size of the breaks $\mathbf{Q}_i$ and the probabilities $\Pr(\kappa_{1ij,\tau} = 1)$ and $\Pr(\kappa_{2i,\tau} = 1)$ for which our priors are quite informative. This is also important because these priors restrict the maximum number of breaks of maximum magnitude and therefore help to identify the factor exposures, which is otherwise rather problematic because linear multifactor models are subject to well-known indeterminacy problems upon rotations of factors and risk premia (see e.g., McCulloch and Rossi, 1991). The prior shape parameters for the probability of breaks in the dynamics of the price sensitivities is set to be $a_{ij} = 3.2$ and $b_{ij} = 60$. As such,

$$E[\pi_{ij}] = \frac{3.2}{3.2 + 60} = 0.05 \qquad \text{and} \qquad Std[\pi_{ij}] = \left(\frac{3.2 \times 60}{(3.2 + 60)^2(3.2 + 60 + 1)}\right)^{1/2} = 0.03$$

which means an expected 5% prior probability of a random shock in the dynamics of factor loadings. With respect to the idiosyncratic volatility, the shape hyperparameters are set to be $a_{i\nu} = 1$ and $b_{i\nu} = 99$, such that

$$E[\pi_{i\nu}] = \frac{1}{1 + 99} = 0.01 \qquad \text{and} \qquad Std[\pi_{i\nu}] = \left(\frac{99}{100^2 \times 101}\right)^{1/2} = 0.01$$

which set the expected prior probability of having a break in the dynamics of idiosyncratic risks to be equal to 1%. These small prior probabilities makes the modelling dynamics more parsimonious, mitigating the magnitude of prior information, letting the data speak about the likelihood of random breaks. The prior beliefs on the size of the breaks are inverse-gamma distributed. The prior scale hyper-parameters $\gamma_{ij}, \gamma_{i\nu}$ and the $\delta_{ij}, \delta_{i\nu}$ degrees of freedom are calibrated supporting a prior view for premiums to be normally distributed with zero mean and variance such that there is 95% probability that annualized premia are smaller in absolute value than the larger between the absolute value between the minimum and the maximum return observed in the sample for all the assets. Finally, the prior residual variance is centered at about 10, a value that appeared in the higher range of the maximum likelihood estimates. All other priors imply that the posteriors tend to be centered around their maximum likelihood estimates which eases comparisons with the existing literature.

[23] Groen, Paap and Ravazzolo (2012) discuss prior sensitivity analysis and MCMC convergence tests, see Appendices C and D, which are also used in this paper. Specific prior values and results for convergence tests are available upon request.

Appendix B: Variance Ratio and 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. (9) decomposes excess asset returns in a component related to risk, represented by the term $\sum_{j=1}^{K} \lambda_{j,t}\beta_{ij,t|t-1}$ plus a residual $\lambda_{0,t} + e_{i,t}$. In principle, a multi-factor model is as good as the implied percentage of total variation in excess returns explained by its first component, $\sum_{j=1}^{K} \lambda_{j,t}\beta_{ij,t|t-1}$. However, here we should recall that even though (9) refers to excess returns, it remains a statistical implementation of the framework in (1). This implies that in practice it may be naive to expect that $\sum_{j=1}^{K} \lambda_{j,t}\beta_{ij,t|t-1}$ be able to explain much of the variability in excess returns. A more sensible goal seems to be that $\sum_{j=1}^{K} \lambda_{j,t}\beta_{ij,t|t-1}$ 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 $t-1$, $\mathbf{Z}_{t-1}$,

$$x_{i,t} = \delta_{i0} + \sum_{m=1}^{M} \delta_{im} Z_{m,t-1} + \xi_{i,t}, \tag{26}$$

to compute the sample variance of fitted values,

$$Var[P(x_{it}|\mathbf{Z}_{t-1})] \equiv Var\left[\widehat{\delta}_{i0} + \sum_{m=1}^{M} \widehat{\delta}_{im} Z_{m,t-1}\right], \tag{27}$$

where the notation $P(x_{it}|\mathbf{Z}_{t-1})$ means "linear projection" of x_{it} on a set of instruments, $\mathbf{Z}_{t-1}$. Second, for each asset $i = 1, ..., N$, a time series of fitted (posterior) risk compensations, $\sum_{j=1}^{K} \lambda_{j,t}\beta_{ij,t|t-1}$, is regressed onto the instrumental variables,

$$\sum_{j=1}^{K} \lambda_{j,t}\beta_{ij,t|t-1} = \widehat{\delta}'_{i0} + \sum_{m=1}^{M} \widehat{\delta}'_{im} Z_{m,t-1} + \xi'_{i,t} \tag{28}$$

to compute the sample variance of fitted risk compensations:

$$Var\left[P\left(\sum_{j=1}^{K} \lambda_{j,t}\beta_{ij,t|t-1}|\mathbf{Z}_{t-1}\right)\right] \equiv Var\left[\widehat{\delta}'_{i0} + \sum_{m=1}^{M} \widehat{\delta}'_{im} Z_{m,t-1}\right]. \tag{29}$$

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

$$VR1 \equiv \frac{Var\left[P\left(\sum_{j=1}^{K}\lambda_{j,t}\beta_{ij,t|t-1}|\mathbf{Z}_{t-1}\right)\right]}{Var[P(x_{it}|\mathbf{Z}_{t-1})]} > 0 \tag{30}$$

$$VR2 \equiv \frac{Var\left[P\left(x_{i,t}-\sum_{j=1}^{K}\lambda_{j,t}\beta_{ij,t|t-1}|\mathbf{Z}_{t-1}\right)\right]}{Var[P(x_{it}|\mathbf{Z}_{t-1})]} > 0. \tag{31}$$

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 our B-TVB-SV framework, drawing from the joint posterior densities of the factor loadings $\beta_{ij,t|t-1}$ and the implied risk premia $\lambda_{j,t}$, $i=1,...,N$, $j=1,...,K$, and $t=1,...,T$, and holding the instruments fixed over time, it is possible to compute VR1 and VR2 in correspondence to each of such draws and hence obtain their posterior distributions.[24]

Finally, the predictable variation of returns due to the multi-factor model may be further decomposed into the components imputed to each of the individual systematic risk factors, by computing the factoring of $Var[P(\sum_{j=1}^{K}\lambda_{j,t}\beta_{ij,t|t-1}|\mathbf{Z}_{t-1})]$ as

$$\sum_{j=1}^{K} Var\left[P\left(\lambda_{j,t}\beta_{ij,t|t-1}|\mathbf{Z}_{t-1}\right)\right] + \sum_{j=1}^{K}\sum_{k=1}^{K} Cov[P\left(\lambda_{j,t}\beta_{ij,t|t-1}|\mathbf{Z}_{t-1}\right), P\left(\lambda_{k,t}\beta_{ik,t|t-1}|\mathbf{Z}_{t-1}\right)] \tag{32}$$

and tabulating $Var\left[P\left(\lambda_{j,t}\beta_{ij,t|t-1}|\mathbf{Z}_{t-1}\right)\right]$ for $j=1,...,K$ as well as the residual factor $\sum_{j=1}^{K}\sum_{k=1}^{K}$ $Cov[P\left(\lambda_{j,t}\beta_{ij,t|t-1}|\mathbf{Z}_{t-1}\right), P\left(\lambda_{k,t}\beta_{ik,t|t-1}|\mathbf{Z}_{t-1}\right)]$ to pick up any interaction terms. Note that because of the existence of the latter term, the equality

$$\sum_{j=1}^{K}\frac{Var\left[P\left(\lambda_{j,t}\beta_{ij,t|t-1}|\mathbf{Z}_{t-1}\right)\right]}{Var\left[P\left(\sum_{j=1}^{K}\lambda_{j,t}\beta_{ij,t|t-1}|\mathbf{Z}_{t-1}\right)\right]} = 1 \tag{33}$$

fails to hold, i.e., the sum of the K risk compensations should not equal the total predictable variation from the asset pricing model because of the covariance among individual risk compensations. This derives from the fact that even though in (1) the risk factors are assumed to be orthogonal, this does not imply that their time-varying total risk compensations ($\lambda_{j,t}\beta_{ij,t|t-1}$ for $j=1,...,K$) should be orthogonal.

[24]Notice that $VR1=1$ does not imply that $VR2=0$ and viceversa, because

$$Var[P(x_{it}|\mathbf{Z}_{t-1})] \neq Var\left[P\left(\sum_{j=1}^{K}\hat{\lambda}_{j,t}\hat{\beta}_{ij,t|t-1}|\mathbf{Z}_{t-1}\right)\right] + Var\left[P\left(r_{i,t}-\hat{\theta}_{i0}-\sum_{m=1}^{M}\hat{\theta}_{im}Z_{m,t-1}|\mathbf{Z}_{t-1}\right)\right].$$

Table 1: Descriptive Statistics

This table reports the descriptive statistics for each of the 23 portfolios used in the empirical analysis as well as the risk factors and the instrumental variables. Data are monthly and cover the sample period 1972:01 - 2011:12.

Portfolio/Factor	**Mean**	**Median**	**Std. Dev.**	**Sharpe Ratio**
	10 Industry Portfolios, Value-Weighted			
Non-Durable Goods	1.107	1.135	4.473	0.248
Durable Goods	0.809	0.815	6.649	0.122
Manufacturing	0.988	1.195	5.195	0.190
Energy	1.163	0.990	5.672	0.205
High-Tech	0.924	0.950	6.897	0.134
Telecommunications	0.948	1.175	4.891	0.194
Shops and Retail	0.974	1.060	5.447	0.179
Healthcare	0.990	1.050	5.094	0.194
Utilities	0.933	0.995	4.142	0.225
Other	0.871	1.320	5.439	0.160
	10 Size-Sorted Portfolios, Value-Weighted			
Decile 1	1.073	1.205	6.347	0.169
Decile 2	1.083	1.390	6.491	0.167
Decile 3	1.125	1.545	6.162	0.182
Decile 4	1.089	1.500	5.952	0.183
Decile 5	1.127	1.680	5.811	0.194
Decile 6	1.081	1.180	5.412	0.200
Decile 7	1.088	1.255	5.382	0.202
Decile 8	1.024	1.275	5.262	0.195
Decile 9	0.986	1.335	4.853	0.203
Decile 10	0.844	1.075	4.473	0.189
	Bond Returns			
10-Year T-Note	0.679	0.628	2.299	0.295
5-Year T-Note	0.635	0.585	1.629	0.390
Baa Corp. Bond (10-20 years)	0.831	0.863	3.237	0.257
	Economic Risk Factors			
Excess Value-Weighted Mkt	0.452	0.800	4.681	0.097
Default Premium	0.192	0.461	3.481	
Term Spread	0.000	0.000	0.406	
Industrial Prod. Growth	0.186	0.256	0.755	
Real Per-capita Cons. Growth	0.255	0.262	0.338	
Real T-Bill Interest Rate	0.087	0.102	0.357	
Unexpected Inflation	0.000	-0.016	0.301	
Bond Risk Factor	1.093	0.982	1.944	0.562
Liquidity Factor	0.497	0.232	3.621	0.137
	Instrumental Variables			
Term Yield Spread	1.715	1.910	1.329	
Credit Yield Spread	1.111	0.960	0.488	
Dividend Yield	3.029	2.952	1.259	

Table 2: Risk Premia

This table reports statistics describing the posterior distribution of the risk premia on each factor across different model specifications. Data are monthly and cover the sample period 1972:01 - 2011:12. The first ten years of monthly data are used to calibrate the priors for all the models except for the standard two-step Fama-MacBeth procedure.

Full sample (Jan. 1982 - Dec. 2011)							
	Average	**Std. Error**	**t-stat**	**p-value**	**2.5%**	**50%**	**97.5%**
Two-step Fama-MacBeth approach							
Intercept	0.2909	0.1363	2.1350	**0.0336**	-3.3346	0.3281	3.3471
Market	0.2593	0.1408	1.8414	**0.0665**	-8.7553	0.6739	7.9319
Credit Spread	0.2208	0.2706	0.8161	0.4151	-4.5672	0.3022	4.8480
Term spread	0.0042	0.0347	0.1201	0.9045	-1.0198	-0.0018	1.1440
IP Growth	-0.0130	0.0092	-1.4086	0.1600	-0.3368	-0.0210	0.3218
Real Consuption Growth	0.0061	0.0039	1.5485	0.1226	-0.1917	-0.0006	0.2309
Real T-bill Rate	-0.0264	0.0412	-0.6414	0.5217	-1.4068	0.0145	1.4703
Unexpected Inflation	-0.0085	0.0062	-1.3670	0.1727	-0.2079	-0.0134	0.2100
Bond Risk Factor	-0.4633	0.1883	-2.4598	**0.0145**	-6.9685	-0.3831	5.1747
Liquidity Factor	0.4012	0.3471	1.1558	0.2487	-12.4488	0.1301	12.9321
Bayesian model with time-varying betas and idiosyncratic risk							
Intercept	0.4125	0.2924	1.4108	0.1593	0.3406	0.4913	0.6432
Market	0.3391	0.1298	2.6119	**0.0095**	0.1207	0.3482	0.5515
Credit Spread	-0.1339	0.1145	-1.1688	0.2434	-0.0471	0.1291	0.3172
Term Spread	-0.0149	0.0306	-0.4880	0.6259	-0.0334	0.0144	0.0616
IP Growth	0.0190	0.0076	2.4940	**0.0132**	0.0031	0.0188	0.0231
Real Consuption Growth	0.0020	0.0044	0.4575	0.6476	-0.0054	0.0018	0.0090
Real T-bill Rate	0.0199	0.0300	0.6616	0.5087	-0.0279	0.0187	0.0682
Unexpected Inflation	-0.0206	0.0064	-3.2095	**0.0015**	-0.0211	-0.0148	-0.0007
Bond Risk Factor	-0.0259	0.0719	-0.3605	0.7187	-0.1449	-0.0218	0.0916
Liquidity Factor	0.3172	0.1560	2.0341	**0.0428**	0.0312	0.3214	0.5719
Bayesian time-varying parameter model (with stochastic volatility)							
Intercept	0.5862	0.0787	7.4482	**0.0000**	0.4575	0.5889	0.7172
Market	0.2197	0.0988	2.2237	**0.0269**	0.0472	0.2220	0.3786
Credit Spread	0.0139	0.0919	0.1516	0.8796	-0.1381	0.0100	0.1710
Term Spread	0.0030	0.0213	0.1401	0.8887	-0.0312	0.0020	0.0382
IP Growth	-0.0079	0.0075	-1.0473	0.2958	-0.0209	-0.0077	0.0036
Real Consuption Growth	0.0047	0.0040	1.1780	0.2397	-0.0013	0.0046	0.0114
Real T-bill Rate	0.0067	0.0216	0.3094	0.7573	-0.0278	0.0055	0.0424
Unexpected Inflation	-0.0092	0.0054	-1.6933	**0.0914**	-0.0183	-0.0090	-0.0001
Bond Risk Factor	-0.0126	0.0492	-0.2550	0.7989	-0.0982	-0.0093	0.0666
Liquidity Factor	0.2071	0.1050	1.9726	**0.0495**	0.0325	0.2066	0.3720
Bayesian model with time-varying betas (No stochastic volatility)							
Intercept	0.5550	0.2775	1.9996	**0.0464**	0.2500	0.5421	0.8540
Market	0.3006	0.1448	2.0758	**0.0388**	0.0291	0.3028	0.5814
Credit Spread	0.1162	0.1582	0.7346	0.4631	-0.1080	0.0967	0.3963
Term Spread	0.0130	0.0550	0.2368	0.8130	-0.0812	0.0105	0.1120
IP Growth	-0.0067	0.0101	-0.6573	0.5115	-0.0218	-0.0060	0.0111
Real Consuption Growth	0.0030	0.0063	0.4788	0.6324	-0.0072	0.0026	0.0141
Real T-bill Rate	0.0191	0.0498	0.3837	0.7014	-0.0620	0.0165	0.1016
Unexpected Inflation	-0.0024	0.0080	-0.3066	0.7594	-0.0172	-0.0021	0.0103
Bond Risk Factor	0.0431	0.1113	0.3847	0.7007	-0.1242	0.0313	0.2512
Liquidity Factor	0.0474	0.2419	0.1958	0.8449	-0.3394	0.0229	0.4512

Table 3: Marginal Likelihoods and Bayes Factors Across Alternative Model Specifications

This table reports the values of the marginal log-likelihoods and the relative Bayes Factors for different model specifications. The values reported are also disaggregated by computing the contributions coming from each of the portfolios under investigation. *B-TVB-SV* stands for Bayesian time-varying betas, stochastic volatility model, while *B-TVB* and *B-TVP* are, respectively, the dynamic Bayesian model restricted to have constant conditional volatility and random-walk betas. *Fama-MacBeth* is the standard two-step procedure. *BF1* is the Bayes Factor for the B-TVB-SV model vs. the no-stochastic volatility restriction. Likewise, *BF2* and *BF3* are the Bayes Factors comparing the B-TVB-SV model with the B-TVP and the Fama-MacBeth approaches, respectively.

	B-TVB-SV	**B-TVB**	**B-TVP**	**Fama-MacBeth**	**BF1**	**BF2**	**BF3**
			10 Industry Portfolios, Value-Weighted				
Non Durable Goods	**-445.39**	-1408.71	-635.40	-3131.83	963.32	190.00	2686.44
Durable Goods	**-700.77**	-1980.33	-832.07	-4412.78	1279.56	131.29	3712.01
Manufacturing	**-330.98**	-1199.96	-522.95	-3851.11	868.98	191.97	3520.13
Energy	**-789.61**	-1793.14	-821.64	-2687.83	1003.53	32.03	1898.22
High Tech	**-571.53**	-1732.31	-717.08	-7269.27	1160.78	145.55	6697.74
Telecommunications	**-614.18**	-1634.38	-734.46	-3353.11	1020.20	120.28	2738.93
Shops and Retail	**-481.33**	-1370.61	-648.98	-4271.70	889.29	167.65	3790.38
Health	**-613.64**	-1591.56	-706.14	-3107.84	977.92	92.50	2494.21
Utilities	**-572.88**	-1684.43	-698.85	-1955.89	1111.55	125.96	1383.01
Other	**-270.90**	-1345.87	-519.80	-6041.50	1074.97	248.90	5770.60
			10 Size-sorted Portfolios, Value-Weighted				
Decile 1	**-620.66**	-1756.62	-725.44	-7211.50	1135.96	104.77	6590.84
Decile 2	**-535.44**	-1632.89	-678.94	-6578.42	1097.46	143.50	6042.98
Decile 3	**-428.11**	-1337.08	-616.02	-6506.86	908.96	187.90	6078.74
Decile 4	**-392.01**	-1262.43	-589.11	-7127.84	870.43	197.11	6735.84
Decile 5	**-335.93**	-1134.01	-543.20	-7517.14	798.08	207.27	7181.21
Decile 6	**-259.32**	-932.83	-506.59	-8008.70	673.51	247.27	7749.38
Decile 7	**-202.41**	-923.42	-468.69	-7121.70	721.01	266.28	6919.29
Decile 8	**-149.14**	-882.35	-446.04	-8924.98	733.21	296.90	8775.84
Decile 9	**-95.261**	-601.39	-365.63	-8158.70	506.13	270.37	8063.44
Decile 10	**-54.063**	-544.78	-329.60	-7820.37	490.72	275.55	7766.31
			Bond Returns				
10 - Yrs Treasury	**-188.98**	-1052.55	-422.69	-9201.70	863.57	233.70	9012.72
5 - Yrs Treasury	**-41.972**	-549.00	-320.55	-7951.90	507.03	278.58	7909.93
Baa Corporate Bonds (10-20 years)	**-185.96**	-1050.90	-420.63	-5522.90	864.94	234.67	5336.94
Overall	**-386.11**	-1278.33	-576.98	-5988.50	892.22	190.87	5602.40

Table 4: Variance Decomposition Tests Across Models

This table reports the results of variance decomposition tests across models. The first two columns show the values from the standard two-step Fama-MacBeth methodology. All rates are in excess of the holding period return on a 1-month T-Bill. VR1 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 $\mathbf{Z}_t$. VR2 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 calculated as in Campbell and Beeler (2012), the lagged yield of a Baa corporate bond, and the lagged spread of long- vs. short-term government bond yields. *B-TVB-SV* stands for Bayesian time-varying betas with stochastic volatility, while *B-TVP* and *B-TVB* are, respectively, Bayesian time-varying parameters and time-varying betas models.

	Fama-MacBeth		B-TVB-SV						B-TVP						B-TVB					
	VR1	VR2	VR1			VR2			VR1			VR2			VR1			VR2		
			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%
10 Industry Portfolios, Value-Weighted																				
Non Durable Goods	0.368	0.695	0.318	0.454	0.530	0.089	0.144	0.294	0.254	0.320	0.450	0.241	0.328	0.418	0.046	0.217	0.749	0.243	0.719	0.835
Durable Goods	0.499	0.450	0.547	0.833	0.969	0.098	0.131	0.213	0.256	0.376	0.789	0.239	0.678	0.783	0.072	0.464	0.802	0.187	0.519	0.832
Manufacturing	0.717	0.205	0.756	0.916	1.053	0.031	0.173	0.309	0.090	0.453	0.816	0.142	0.418	0.934	0.134	0.492	0.818	0.200	0.569	0.845
Energy	0.698	0.354	0.897	0.987	1.090	0.017	0.056	0.067	0.707	0.832	0.952	0.115	0.221	0.305	0.613	0.865	0.938	0.087	0.148	0.296
High Tech	0.660	0.380	0.869	0.941	1.011	0.108	0.138	0.167	0.345	0.546	0.812	0.256	0.431	0.680	0.337	0.614	0.854	0.105	0.309	0.638
Telecommunications	0.494	0.476	0.730	0.888	0.978	0.109	0.154	0.212	0.117	0.436	0.918	0.073	0.636	0.828	0.176	0.418	0.706	0.326	0.533	0.799
Shops and Retail	0.663	0.336	0.615	0.795	0.913	0.045	0.213	0.378	0.112	0.383	0.901	0.139	0.529	0.922	0.190	0.410	0.832	0.159	0.561	0.790
Health	0.428	0.539	0.612	0.830	0.886	0.003	0.116	0.227	0.202	0.479	0.782	0.247	0.518	0.675	0.134	0.490	0.796	0.144	0.560	0.804
Utilities	0.266	0.705	0.432	0.600	0.690	0.017	0.410	0.760	0.066	0.320	0.869	0.134	0.520	0.910	0.101	0.402	0.708	0.266	0.592	0.871
Other	0.278	0.700	0.431	0.615	0.699	0.152	0.357	0.744	0.231	0.375	0.552	0.428	0.592	0.775	0.240	0.385	0.633	0.311	0.618	0.660
10 Size-Sorted Portfolios, Value-Weighted																				
Decile 1	0.314	0.677	0.225	0.309	0.364	0.281	0.648	0.797	0.242	0.342	0.450	0.552	0.681	0.704	0.015	0.296	0.622	0.305	0.750	0.901
Decile 2	0.731	0.182	0.723	0.882	0.985	0.067	0.159	0.255	0.141	0.467	0.842	0.190	0.564	0.761	0.278	0.456	0.689	0.218	0.473	0.704
Decile 3	0.629	0.323	0.611	0.906	0.954	0.002	0.163	0.316	0.278	0.569	0.976	0.131	0.409	0.889	0.168	0.442	0.706	0.187	0.389	0.783
Decile 4	0.603	0.380	0.542	0.820	0.961	0.043	0.178	0.408	0.270	0.520	0.886	0.234	0.516	0.814	0.239	0.488	0.906	0.079	0.476	0.847
Decile 5	0.538	0.497	0.633	0.877	0.999	0.060	0.112	0.311	0.089	0.465	0.951	0.080	0.473	0.820	0.061	0.261	0.552	0.471	0.750	0.906
Decile 6	0.262	0.679	0.620	0.864	1.012	0.030	0.166	0.348	0.070	0.418	0.824	0.120	0.500	0.899	0.154	0.354	0.737	0.198	0.635	0.891
Decile 7	0.448	0.515	0.723	0.846	0.993	0.013	0.183	0.270	0.102	0.478	0.858	0.144	0.360	0.813	0.165	0.372	0.763	0.247	0.604	0.882
Decile 8	0.367	0.650	0.777	0.855	0.942	0.050	0.111	0.276	0.277	0.552	0.834	0.153	0.361	0.737	0.266	0.501	0.809	0.154	0.491	0.693
Decile 9	0.614	0.367	0.686	0.922	1.025	0.053	0.164	0.369	0.169	0.338	0.795	0.095	0.334	0.732	0.059	0.277	0.717	0.244	0.660	0.949
Decile 10	0.585	0.415	0.622	0.768	0.843	0.039	0.229	0.401	0.313	0.599	0.818	0.205	0.460	0.816	0.103	0.354	0.821	0.232	0.616	0.814
Bond Returns																				
10 - Yrs Treasury	0.707	0.249	0.848	0.962	1.099	0.024	0.068	0.112	0.209	0.449	0.837	0.117	0.456	0.759	0.481	0.674	0.808	0.157	0.364	0.482
5 - Yrs Treasury	0.315	0.763	0.819	0.901	1.074	0.047	0.115	0.219	0.307	0.580	0.856	0.199	0.428	0.731	0.282	0.511	0.723	0.246	0.482	0.880
Baa Corp Bonds (10-20 years)	0.443	0.458	0.260	0.376	0.424	0.452	0.563	0.665	0.351	0.467	0.710	0.279	0.512	0.758	0.056	0.245	0.597	0.399	0.675	0.818

Table 5: Average Pricing Errors

This table reports the average pricing errors for each of the models under investigation across different subsamples as well as in the full sample. *B-TVB-SV* stands for Bayesian time-varying betas, stochastic volatility model, while *B-TVB* and *B-TVP* are, respectively, the dynamic Bayesian model restricted to have constant conditional volatility and random-walk betas. *Fama-MacBeth* is the standard two-step procedure. The table reports the average (over time), the posterior standard deviation as well as the confidence interval at the 95% level.

	Average Pricing Errors				
	Mean %	**Std %**	**2.5 %**	**50 %**	**97.5 %**
Panel A: Full-Sample					
B-TVB-SV	**0.2108**	0.0623	0.1363	**0.1902**	0.3231
B-TVB	0.4126	0.1588	0.2512	0.3459	0.7325
B-TVP	0.5401	0.1804	0.3113	0.5061	0.8126
Fama-MacBeth	0.6303	0.0159	0.6107	0.6258	0.6633
Panel B: 1982:01 - 1999:01					
B-TVB-SV	**0.1926**	0.0431	0.1443	**0.1851**	0.2574
B-TVB	0.3935	0.0314	0.3511	0.3921	0.4521
B-TVP	0.5233	0.1653	0.3202	0.4759	0.8101
Fama-MacBeth	0.6278	0.0151	0.6092	0.6238	0.6585
Panel C: 1999:01 - 2011:11					
B-TVB-SV	**0.2624**	0.0672	0.1454	**0.2707**	0.3525
B-TVB	0.4682	0.119	0.2806	0.4501	0.7068
B-TVP	0.6359	0.1993	0.3544	0.6456	0.9027
Fama-MacBeth	0.6354	0.0162	0.6168	0.6321	0.6653
Panel D: 2007:01 - 2011:11					
B-TVB-SV	**0.2891**	0.0613	0.1673	**0.2977**	0.3952
B-TVB	0.5865	0.0906	0.4718	0.5713	0.7559
B-TVP	0.6397	0.1431	0.4029	0.6616	0.8168
Fama-MacBeth	0.6523	0.0179	0.6164	0.6439	0.6799

Figure 1: Mean Posterior Probability of Breaks in Factor Loadings Across Assets/Portfolios

This figure reports the average over the sample of median posterior probabilities of a break in betas across portfolios and factors for both the B-TVB-SV and B-TVB models. The sample period is 1972:01 - 2011:12. The first 120 monthly observations are used as a training sample in order to calibrate the prior distribution for both latent states and parameters. The heating map is reported on the right-hand side.

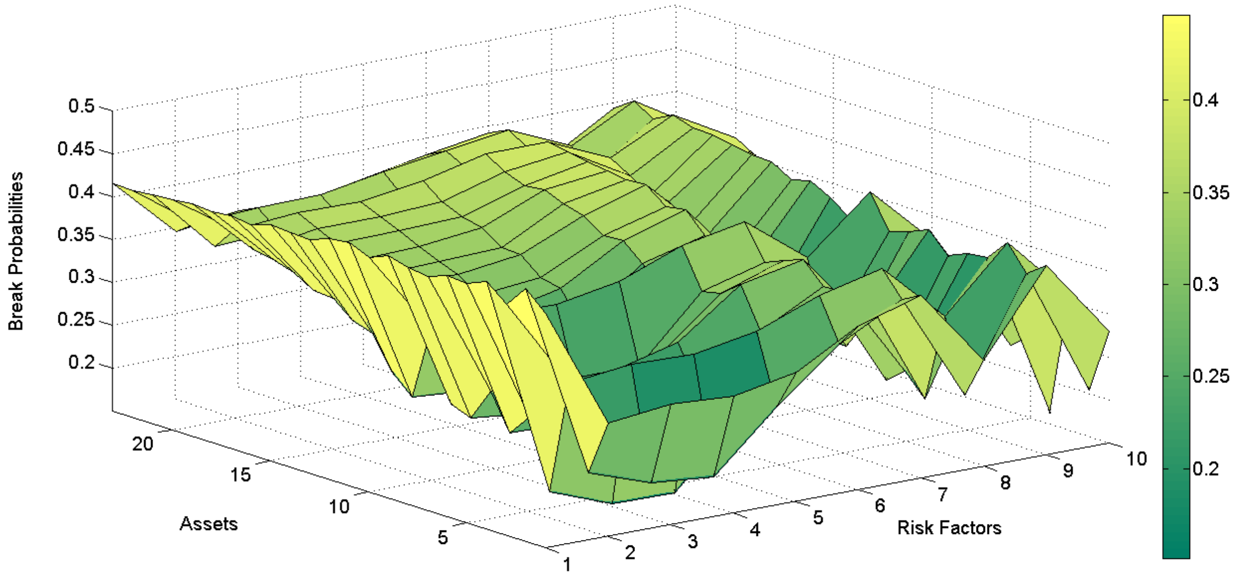

(a) Bayesian Dynamic Model with Instability in Betas and Conditional Volatility

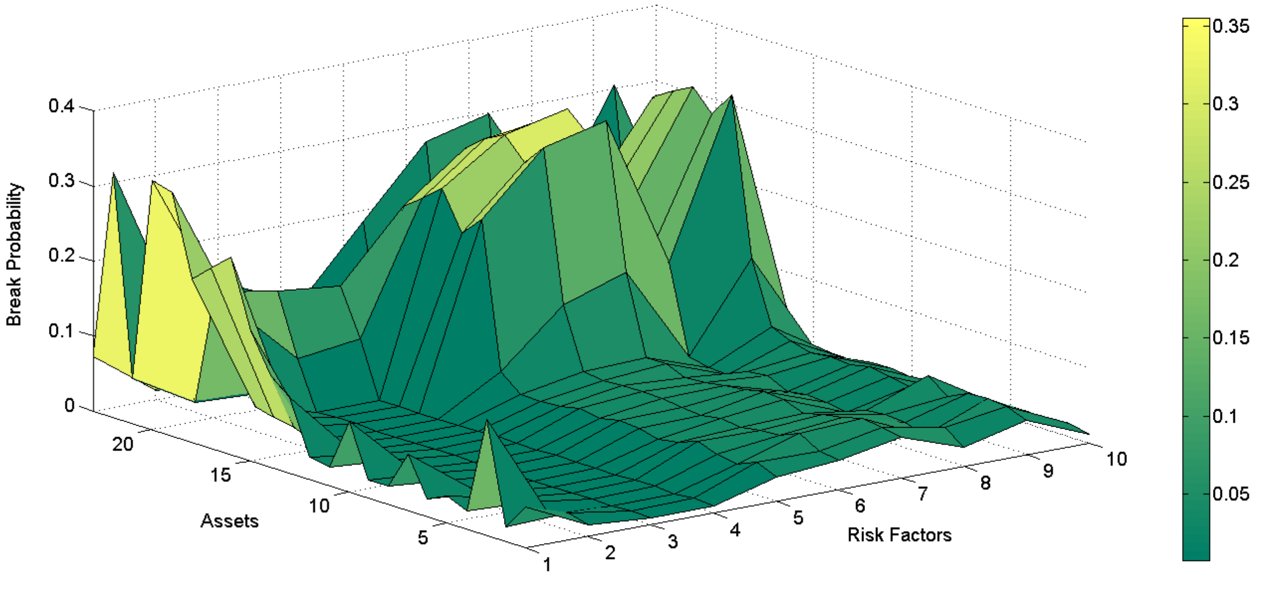

(b) Bayesian Dynamic model with Constant Conditional Volatility

Figure 2: B-TVB-SV Factor Loadings: VW Market Portfolio

This figure reports the time series of the posterior median loadings for the market risk factor estimated from a dynamic Bayesian model with time-varying betas and idiosyncratic risk. The sample period is 1972:01 - 2011:12. The first 120 monthly observations are used as a training sample in order to calibrate the prior distribution for both latent states and parameters. The gray area surrounding posterior median plots represents 95% confidence intervals.

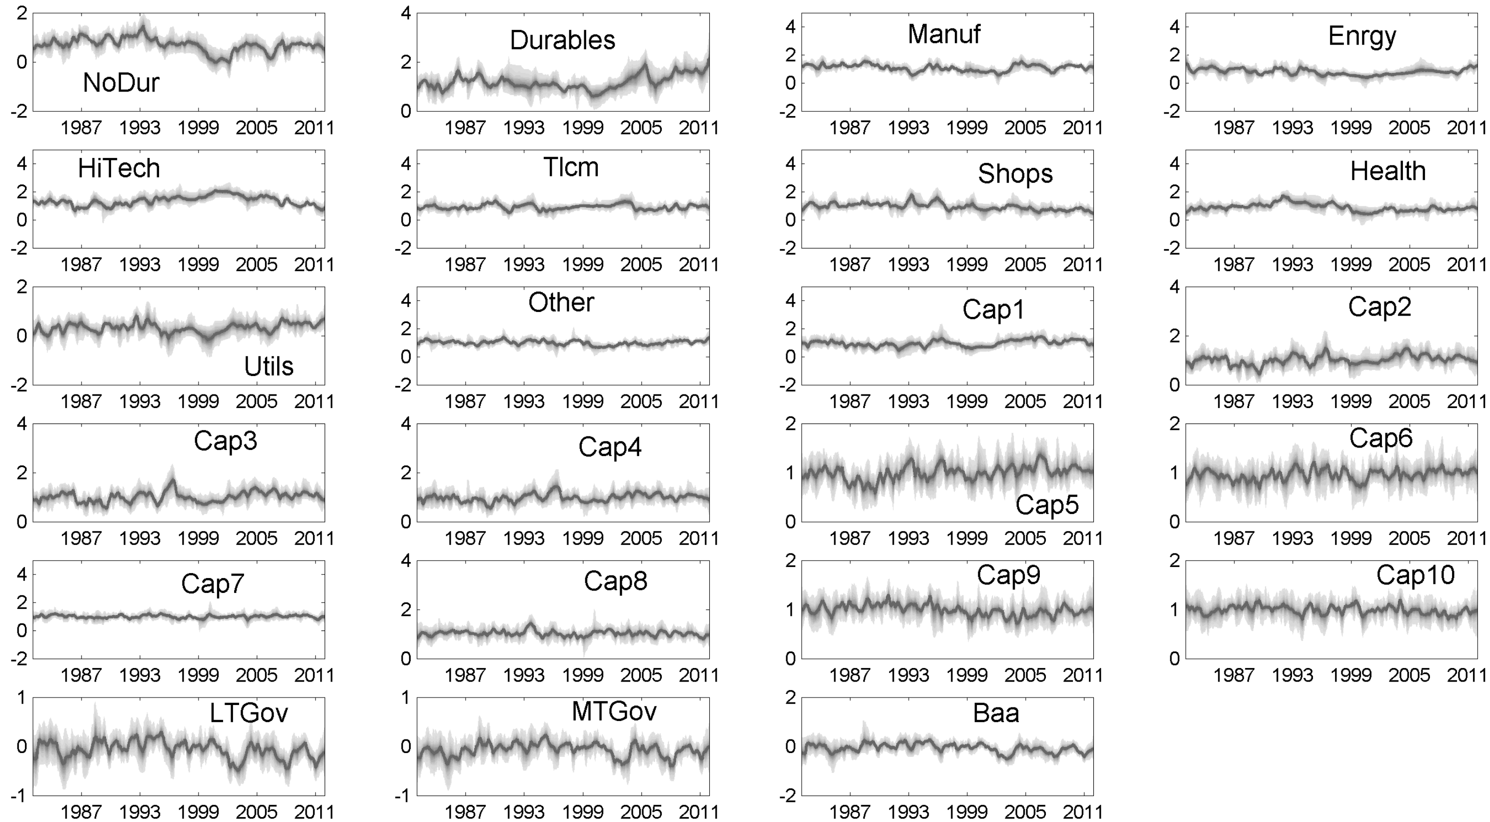

Figure 3: Factor Loadings Estimated with the Fama-MacBeth Approach: VW Market Portfolio

This figure reports the time series of the of the posterior mean loadings for the market risk factor estimated from a naive 5-year rolling-window estimation approach. The sample period is 1972:01 - 2011:12. The red, dashed lines surrounding posterior mean plots represent 95% frequentist confidence intervals. Asymptotic standard errors are computed assuming absence of cross-sectional dependence among the betas estimates.

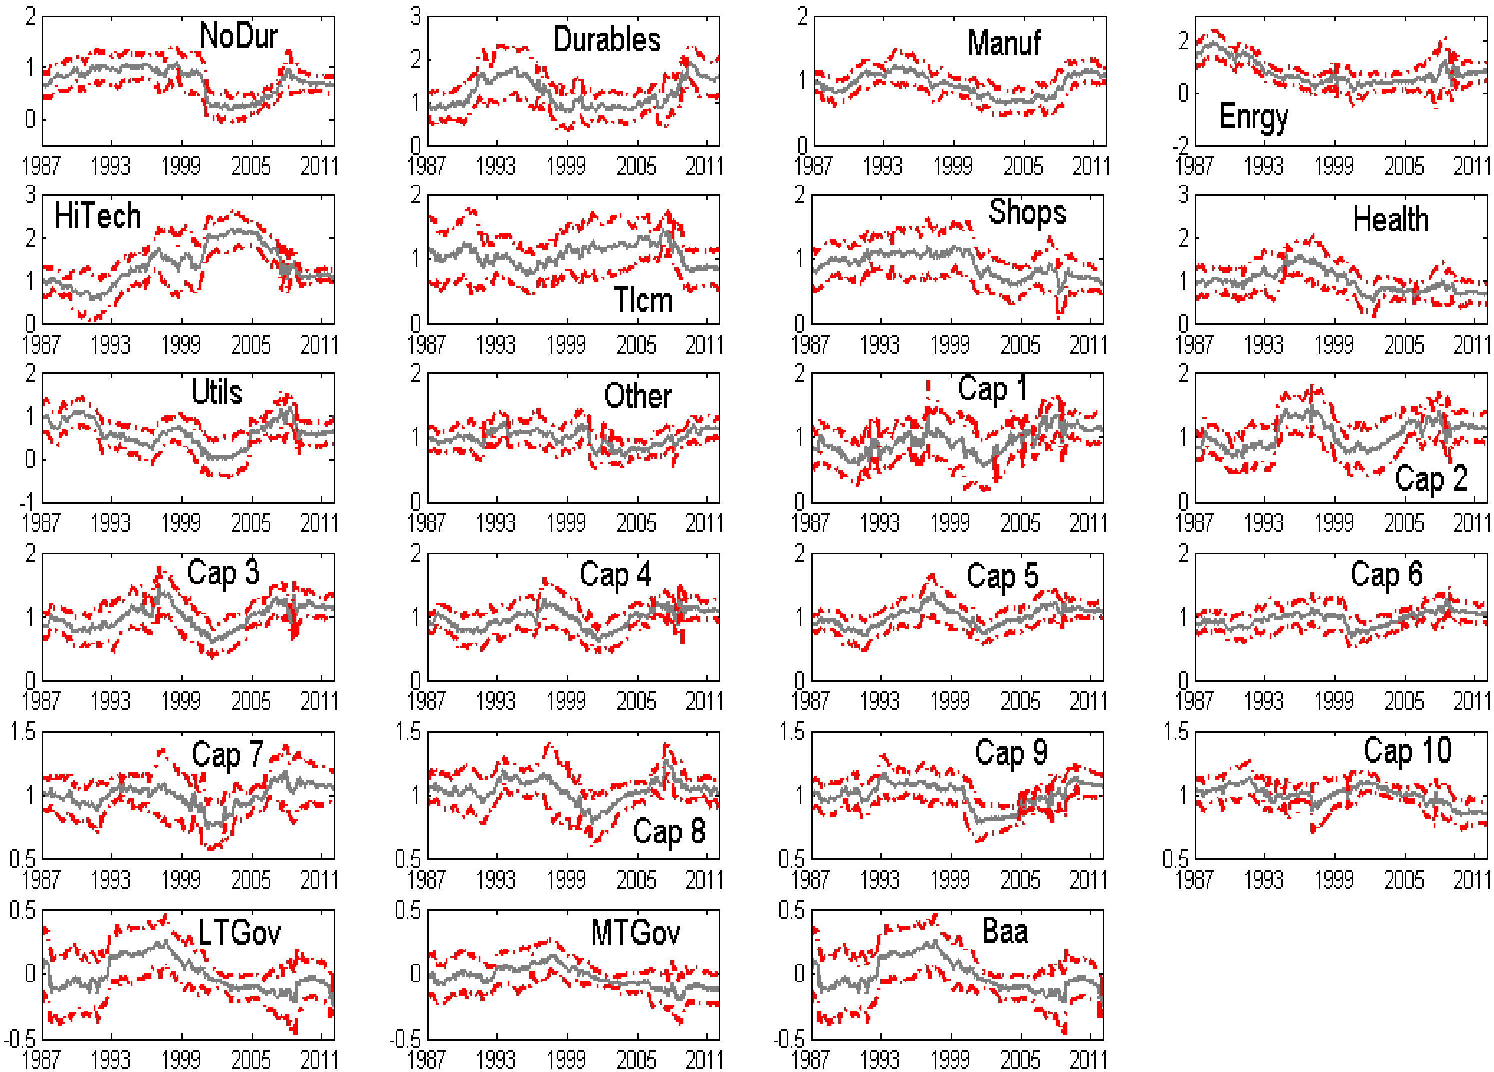

Figure 4: B-TVB-SV Factor Loadings: Term Spread

This figure reports the time series of the posterior median loadings for the term spread risk factor estimated from a dynamic Bayesian model with time-varying betas and idiosyncratic risk. The sample period is 1972:01 - 2011:12. The first 120 monthly observations are used as a training sample in order to calibrate the prior distribution for both latent states and parameters. The gray area surrounding posterior median plots represents 95% confidence intervals.

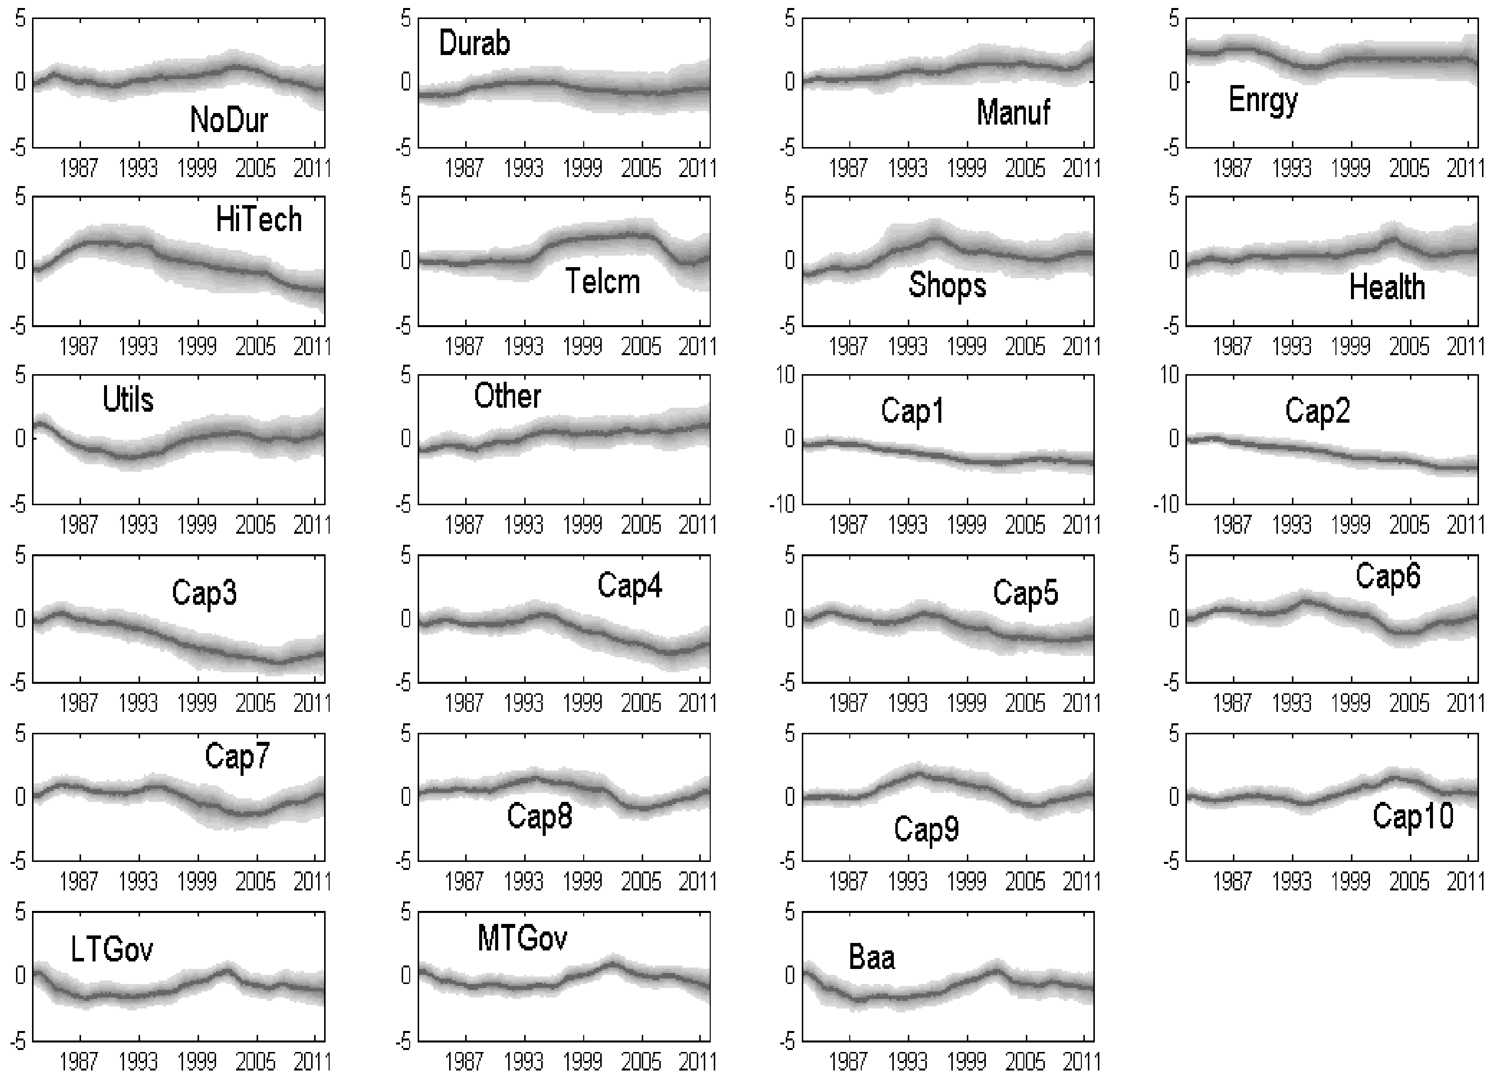

Figure 5: B-TVB-SV Factor Loadings: Industrial Production

This figure reports the time series of the posterior median loadings for the industrial production growth factor estimated from a dynamic Bayesian model with time-varying betas and idiosyncratic risk. The sample period is 1972:01 - 2011:12. The first 120 monthly observations are used as a training sample in order to calibrate the prior distribution for both latent states and parameters. The gray area surrounding posterior median plots represents 95% confidence intervals.

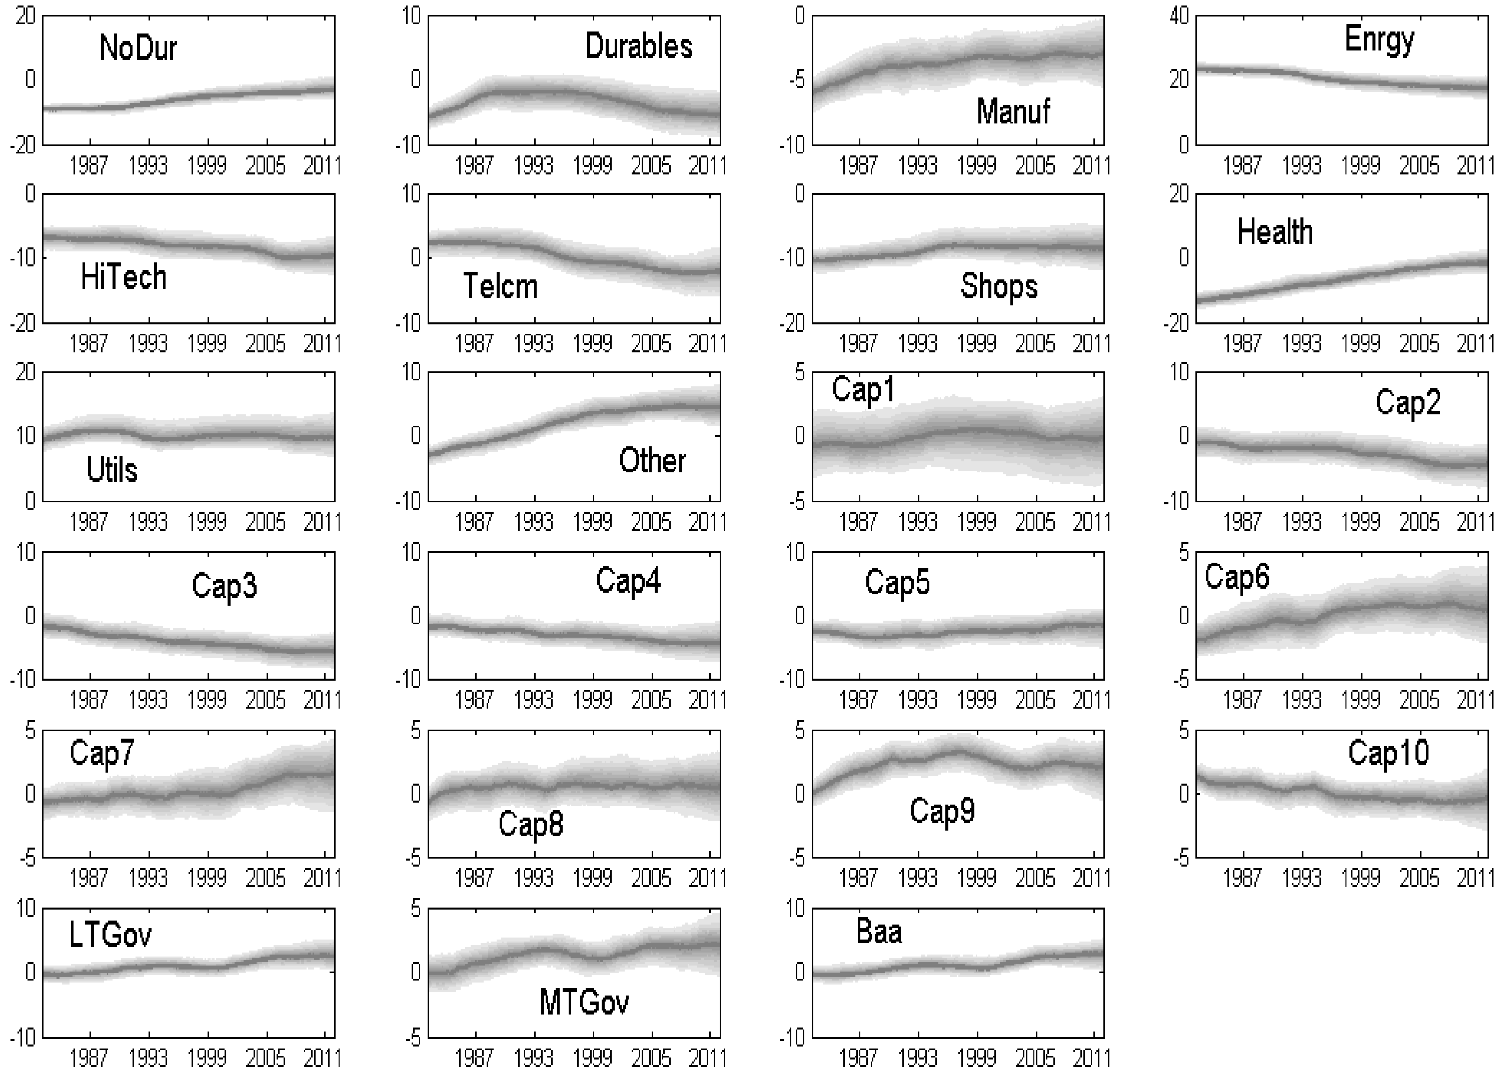

Figure 6: B-TVB-SV Factor Loadings: Unexpected Inflation

This figure reports the time series of the posterior median loadings for the unexpected inflation factor estimated from a dynamic Bayesian model with time-varying betas and idiosyncratic risk. The sample period is 1972:01 - 2011:12. The first 120 monthly observations are used as a training sample in order to calibrate the prior distribution for both latent states and parameters. The gray area surrounding posterior median plots represents 95% confidence intervals.

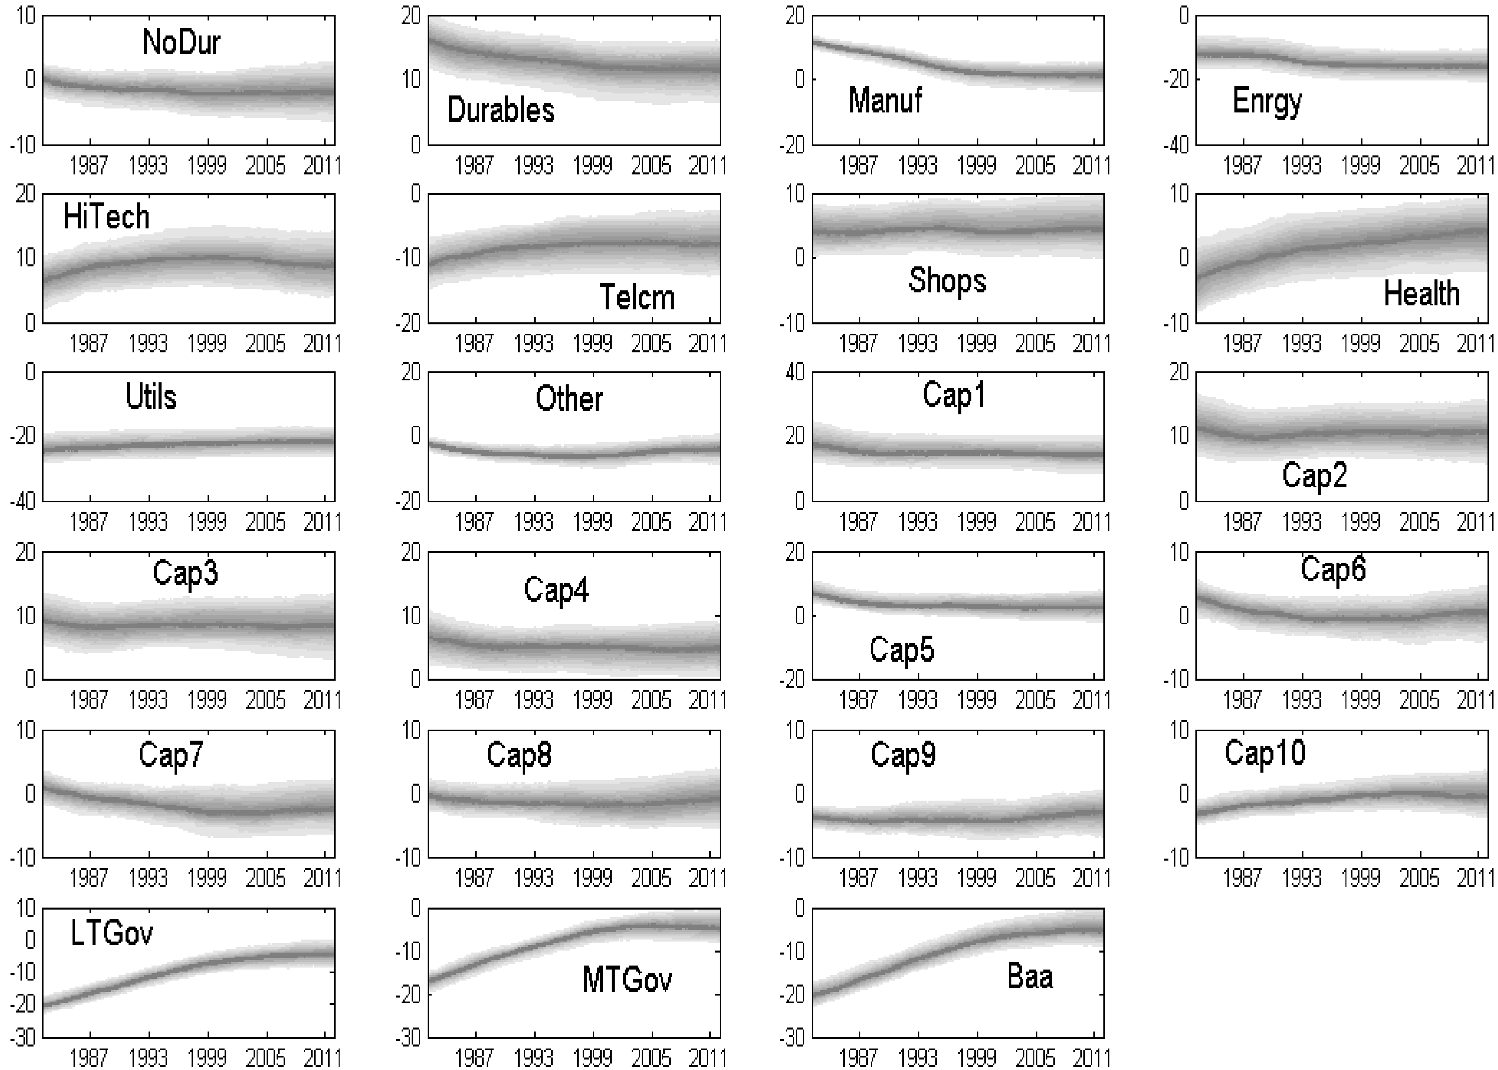

Figure 7: B-TVB-SV Jensen's Alphas

This figure reports the time series of the posterior medians of the Jensen's alpha from a dynamic Bayesian model with time-varying betas and idiosyncratic risk. The sample period is 1972:01 - 2011:12. The first 120 monthly observations are used as a training sample in order to calibrate the prior distribution for both latent states and parameters. The gray area surrounding posterior median plots represents 95% confidence intervals.

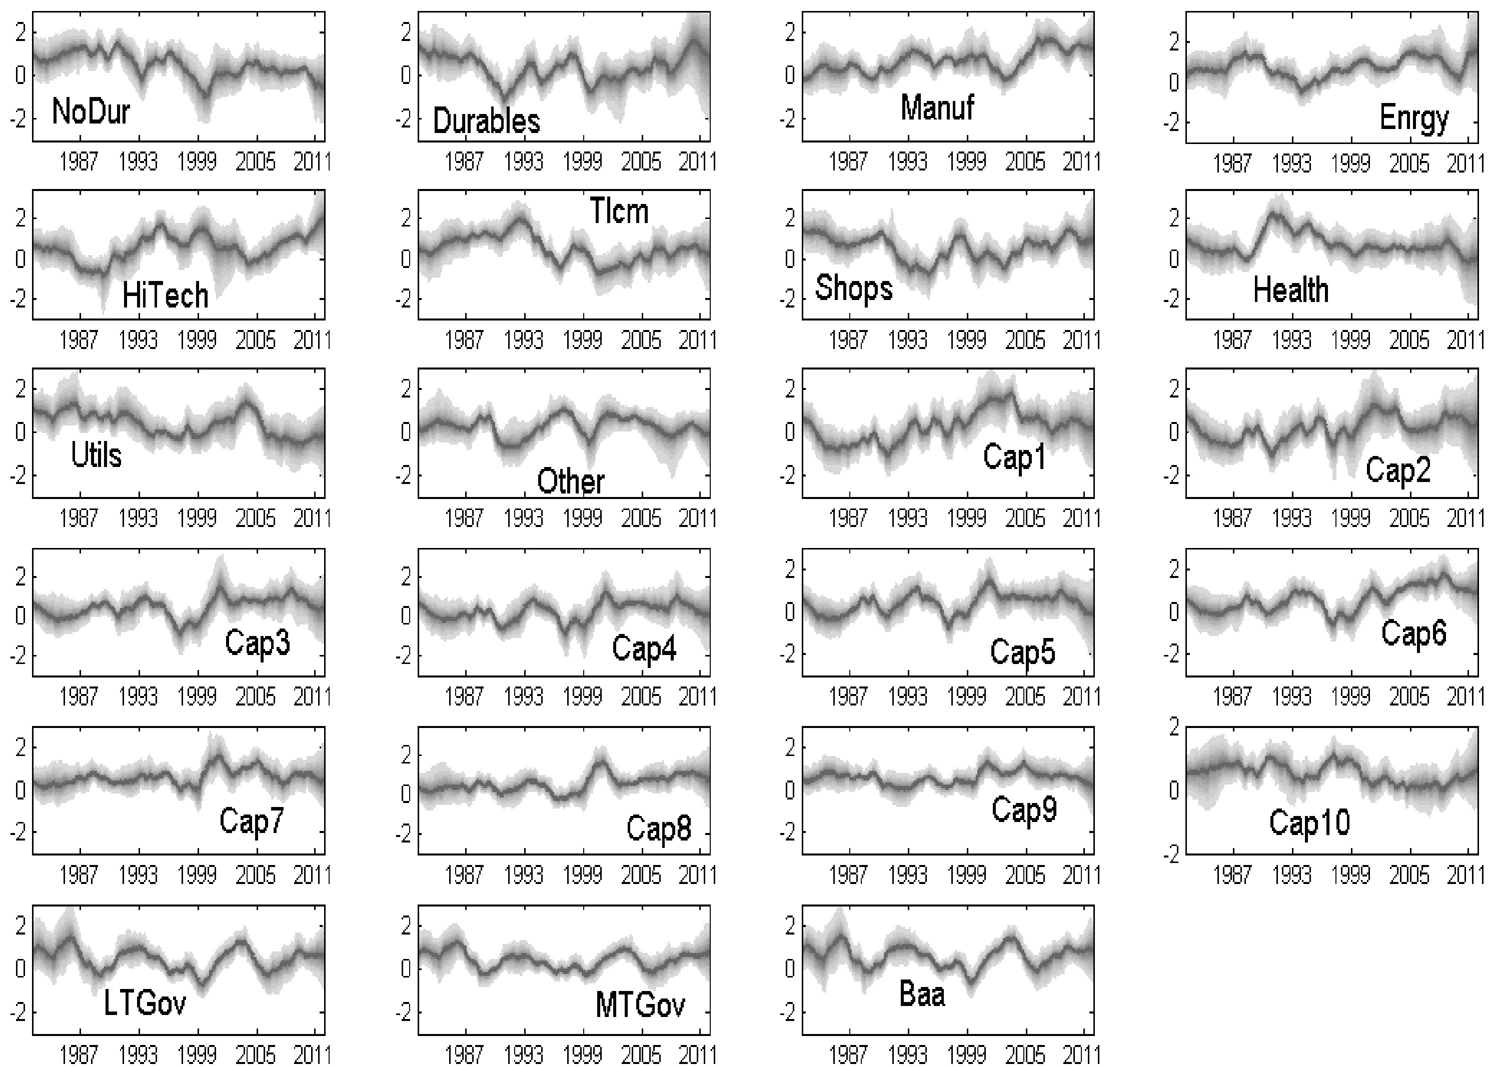

Figure 8: Jensen's Alphas Estimated with the Fama-MacBeth Approach

This figure reports the time series of the of the posterior means of the Jensen's alphas estimated from a naive 5-year rolling-window estimation approach. The sample period is 1972:01 - 2011:12. The red, dashed lines surrounding posterior mean plots represent 95% confidence intervals.

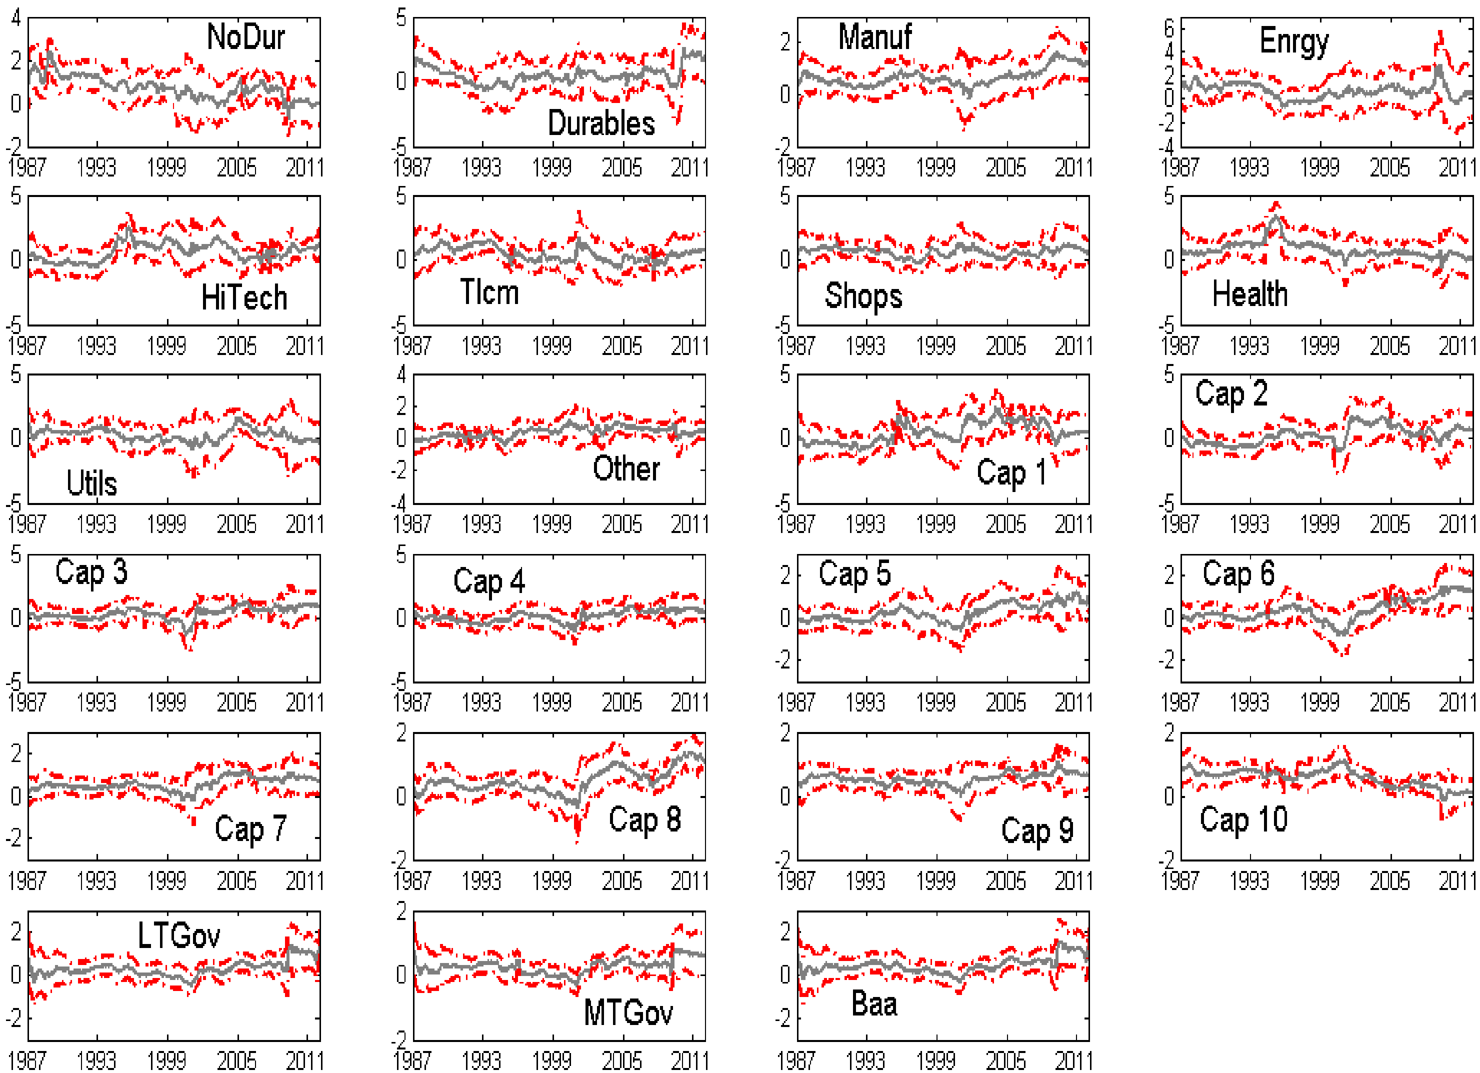

Figure 9: B-TVB-SV Idiosyncratic Risk Dynamics

This figure reports the time series of the posterior medians for idiosyncratic risk estimated from a dynamic Bayesian model with time-varying betas and idiosyncratic risk. The sample period is 1972:01 - 2011:12. The first 120 monthly observations are used as a training sample in order to calibrate the prior distribution for both latent states and parameters. The gray area surrounding posterior median plots represents 95% confidence intervals.

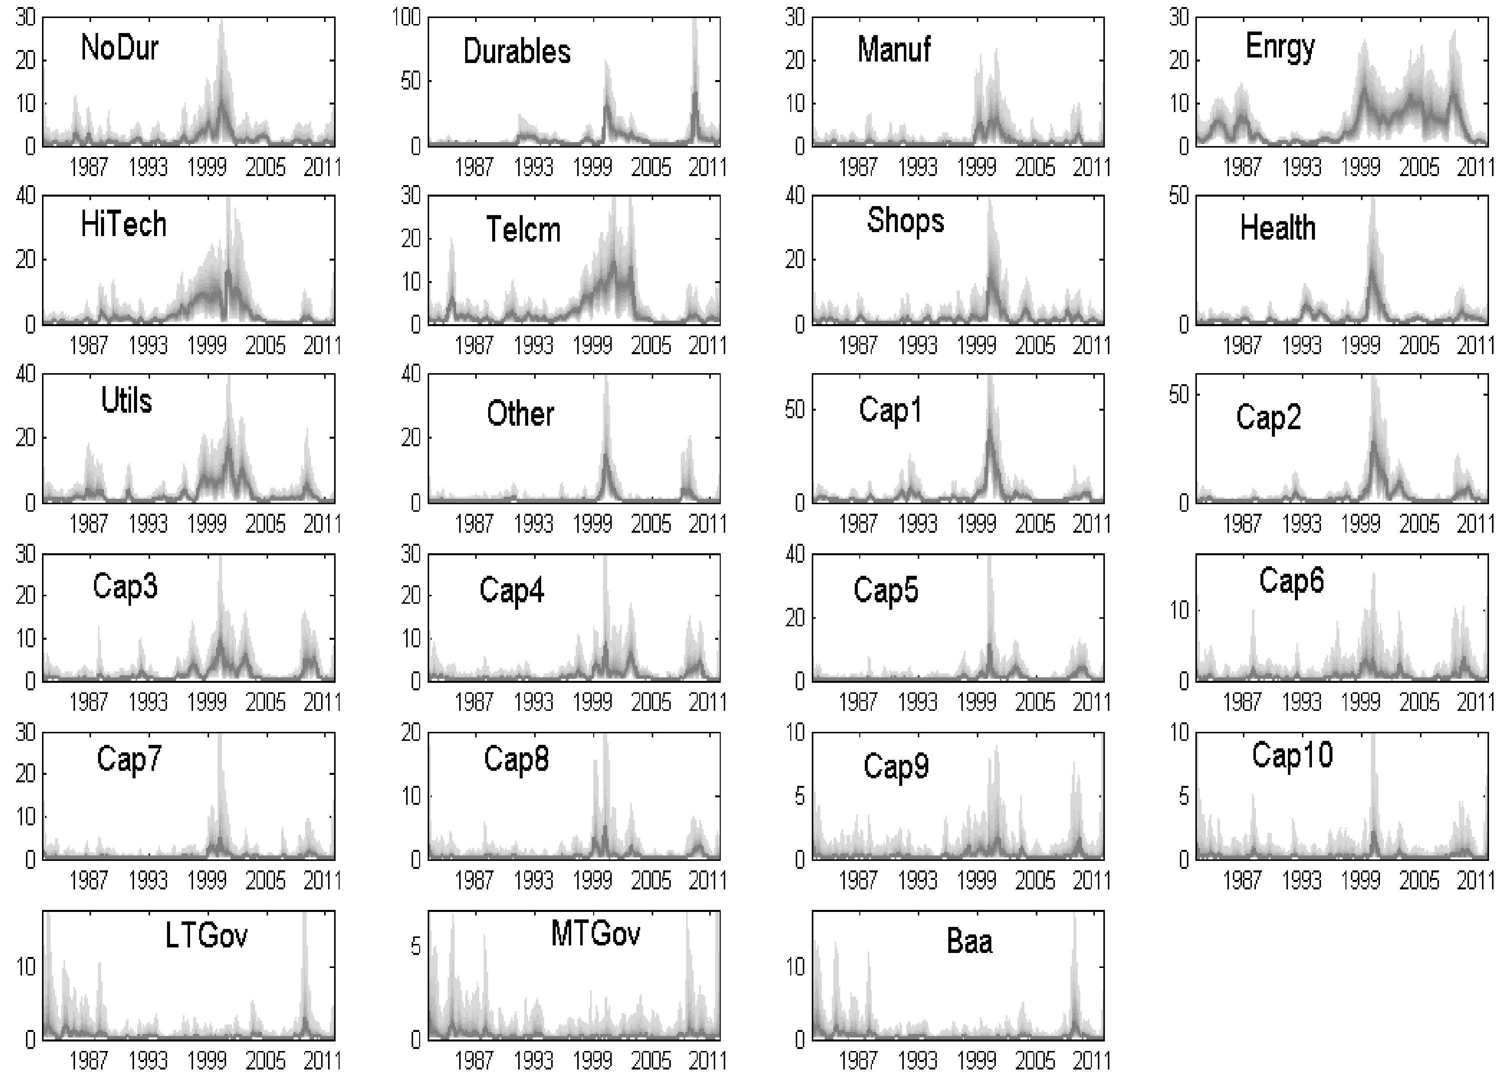

Figure 10: Average Pricing Errors

This figure reports the time series of the average pricing errors. The sample period is 1972:01 - 2011:12. The first 120 monthly observations are used a training sample in order to calibrate the prior distribution for both latent states and parameters. Panel A reports the average pricing error across models. Panel B reports the rescaled values of the average pricing errors for the B-TVB-SV and the B-TVP models, respectively.

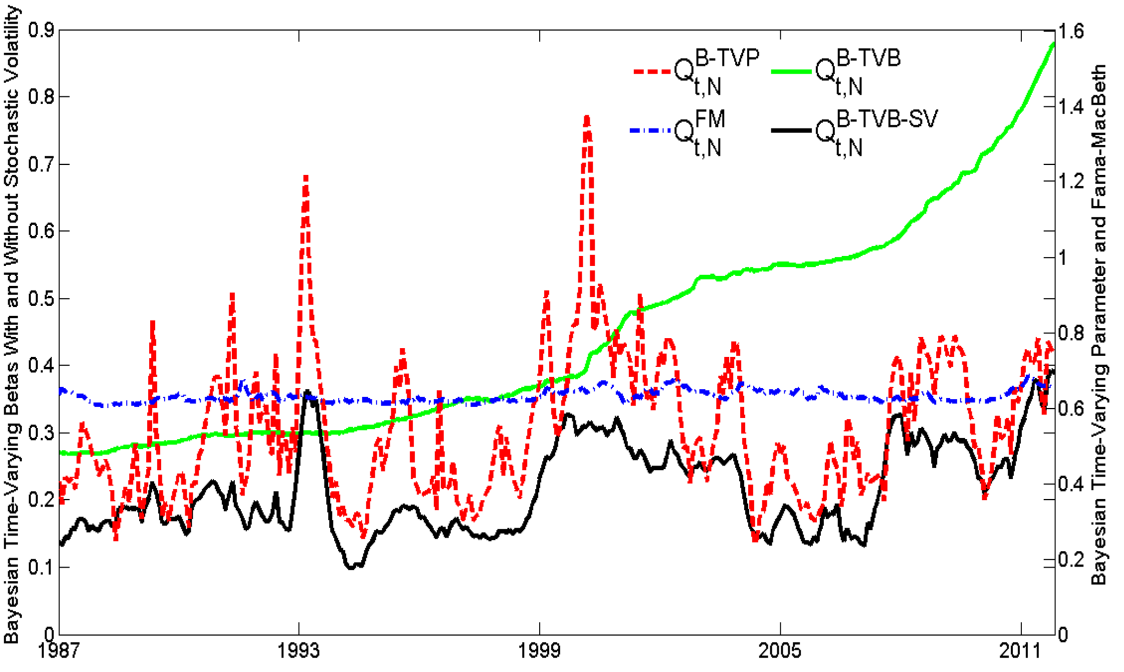

(a) Average Pricing Errors Across Models

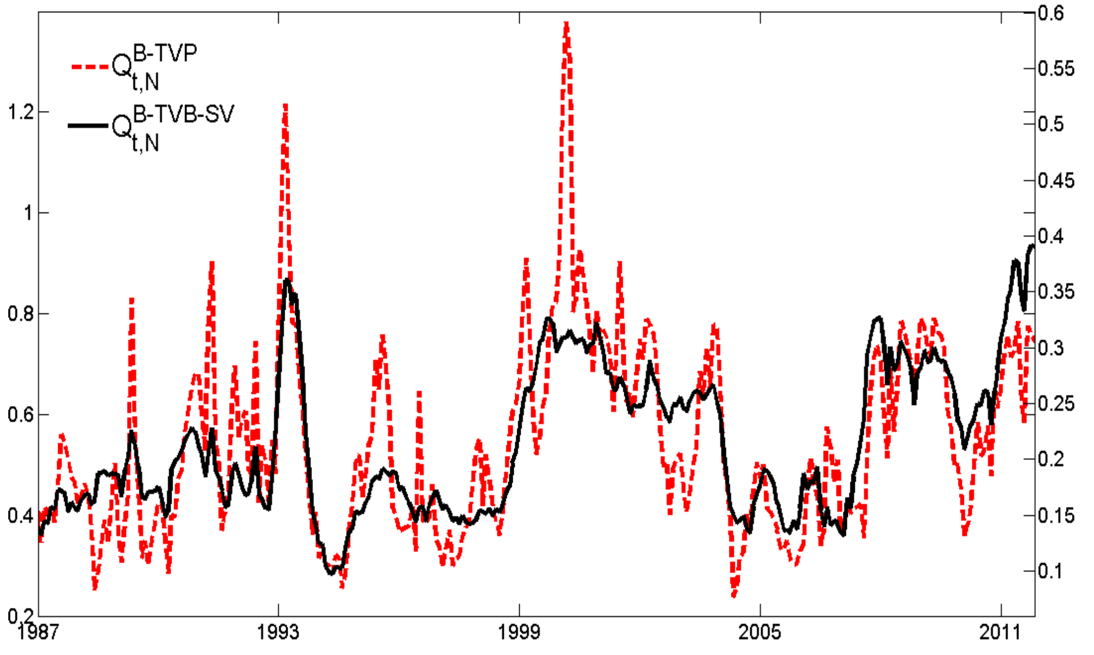

(b) Average Pricing Errors for B-TVB-SV and B-TVP (Rescaled)