



Alternative econometric implementations of multi-factor models of the U.S. financial markets[☆]

Massimo Guidolin^{a,b,*}, Francesco Ravazzolo^c, Andrea Donato Tortora^d

^a Bocconi University, IGIER, Italy

^b CAIR, Manchester Business School, United Kingdom

^c Norges Bank and BI Norwegian Business School, Norway

^d Bocconi University, Milan, Italy

ARTICLE INFO

Article history:

Received 30 January 2011

Received in revised form 19 August 2012

Accepted 13 January 2013

Available online 9 February 2013

JEL classification:

G11

C53

Keywords:

Bayesian estimation

Latent jumps

Stochastic volatility

Linear factor models

ABSTRACT

This paper analyzes the empirical performance of two alternative ways in which multi-factor models with time-varying risk exposures and premia may be estimated. The first method echoes the seminal two-pass approach introduced by Fama and MacBeth (1973). The second approach is based on a Bayesian latent mixture model with breaks in risk exposures and idiosyncratic volatility. Our application to monthly, 1980–2010 U.S. data on stock, bond, and publicly traded real estate returns shows that the classical, two-stage approach that relies on a nonparametric, rolling window estimation of time-varying betas yields results that are unreasonable. There is evidence that most portfolios of stocks, bonds, and REITs have been grossly over-priced. On the contrary, the Bayesian approach yields sensible results and a few factor risk premia are precisely estimated with a plausible sign. Predictive log-likelihood scores indicate that discrete breaks in both risk exposures and variances are required to fit the data.

© 2013 The Board of Trustees of the University of Illinois. Published by Elsevier B.V. All rights reserved.

1. Introduction

Asset pricing is the sub-field of financial economics that investigates the key drivers of the general pricing mechanism (also called the pricing kernel) that underlies observed prices and returns of traded securities (see e.g., Cochrane, 2005, for an introduction). Because most asset pricing models pose non-trivial issues related to their actual estimation (more generally, *implementation*, to also include the choice of securities and asset classes and of the length and frequency of test data), as of lately a related but not less important question has been investigated (see e.g., Singleton, 2006): What are the most appropriate methods to learn about the pricing kernel that underlies the observed cross-section of asset prices

and returns? Our paper offers a contribution to the voluminous literature that has tackled this question specializing to a particular set of models (of the pricing kernel) and applying novel methods to an interesting application. Using 31 years of monthly data on excess returns on 27 key portfolios of securities traded in the U.S., we investigate within a relatively wide class of dynamic latent variable models estimated in a Bayesian state-space framework which models (if any) stand the best chance to implement macro-based linear factor models leading to sensible conclusions concerning the dynamics over time of both factor exposures and risk premia.

The paper has three building blocks. First, we compare two alternative approaches to estimate a standard multifactor asset pricing model (MFAPM) in which the risk factors consists of shocks to observable macroeconomic variables that appear to be commonly tracked by researchers, policy-makers, and the press (e.g., aggregate market returns, the rate of growth of industrial production, changes in the unemployment rate, the spread between long- and short-term nominal rates, etc.). Going back to the seminal work by Chen, Roll, and Ross (1986) there is of course an ever expanding literature that has worked with such a class of models. In particular, Ferson and Harvey (1991) extended the early work on MFAPMs to incorporate the case of time-varying risk premia and exposures. In general terms, a MFAPM has a very simple structure: the risk premium on any asset or portfolio is decomposed as

[☆] We are grateful to seminar participants at Norges Bank, the IREBS 2009 Conference on Real Estate Economics and Finance, and the 4th International Conference on Computational and Financial Econometrics in London (CFE'10) for helpful comments. Andrea Tortora acknowledges financial support from the Marie Curie Early Stage Training Programme. The views expressed in this paper are our own and do not necessarily reflect those of Norges Bank.

* Corresponding author.

E-mail addresses: massimo.guidolin@unibocconi.it, max.guidolin@gmail.com (M. Guidolin), francesco.ravazzolo@Norges-Bank.no (F. Ravazzolo), andrea.tortora@phd.unibocconi.it (A.D. Tortora).

the sum of a number of products between risk exposures (also called betas) to each of the factors and the associated unit price of the factor. Assuming correct specification, the difference at each point in time between actual, realized excess returns and the risk premium implied by the model is called residual or idiosyncratic risk and it is supposed to pick up all the variation in excess returns that is specific to individual portfolios.

Second, our paper jointly uses data on publicly traded stock, bond, and real estate securities (or traded funds invested in the underlying securities), instead of focussing on only one of these asset classes. Therefore, our paper relates to a vast literature that has examined the empirical performance of MFAPMs across asset classes. For instance, Chan, Henderhott, and Sanders (1990) have shown that MFAPMs that include predetermined macroeconomic factors explain a significant proportion of the variation in equity real estate investment trusts (henceforth REITs) returns. Karolyi and Sanders (1998) have extended this evidence and allowed for time-varying risk premia and betas.

Third, our key contribution consists of comparing the heterogeneous results derived from a number alternative implementations of a standard Ferson and Harvey (1991)-style MFAPM. The first approach follows the now standard, Fama and MacBeth (1973) two-stage methodology, first proposed for the plain vanilla CAPM but then extended to the wider class of linear factor models. Fama–MacBeth’s approach uses a first set of rolling window, time series regressions to obtain least-square estimates of the risk exposures, followed by a second-pass set of cross-sectional (across assets or portfolios) regressions that—using the first-pass rolling window betas as inputs—derives time-varying estimates of the associated risk premia. The problems with this methodology are notorious: 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 clearly clashes with them being least squares estimates (hence, sample statistics). Obviously, unless additional assumptions are introduced, this creates a potentially enormous problem with generated regressors being used in the second-step, which tends to make invalid most the inferential statements commonly made when the resulting error-in-variables problems are ignored. Fama–MacBeth’s approach also suffers from another problem: although now common, identifying time-variation in risk exposures and risk premia with a need to perform rolling window least square estimation is surely robust (because, in a way, nonparametric) but always arbitrary and often unsatisfactory.

As a result of these drawbacks of Fama–MacBeth’s approach, we follow a different path based on two pillars: first, time variation in risk exposures, premia, as well as idiosyncratic variances are explicitly modeled as a latent break-point process; 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. Second, the model is estimated using a Bayesian approach that not only is numerically practical but also allows a researcher to feed her own priors in the estimation problem. In the most encompassing of our latent mixture implementations, we model both factor sensitivities and idiosyncratic volatility as latent stochastic processes within a Bayesian framework by means of the mixture innovation approach as in Giordani and Kohn (2008). However, we also consider simpler models nested within the baseline mixture model in which either the breaks are continuous (in the sense that however small, they occur at each point in time, i.e., in a time-varying parameter model) or in which idiosyncratic variances are constant in time.

Our main results can be summarized as follows. The key finding is that—at least in our application and with our data—a standard two-stage Fama–MacBeth approach yields unreasonable economic

implications and is rejected in a statistical sense. First, all of the Bayesian estimates of the loadings (the betas) are considerably smoother than the rolling window, Fama–MacBeth ones, which are instead subject to massive instability, often impossible to interpret. This is a seemingly counter-intuitive result: even though a Bayesian model with latent breaks formally allows the risk exposures to be subject to jumps over time, the resulting posterior densities are actually smoother than what one could retrieve using a naïve rolling window estimation procedure.

Second, the two-step Fama–MacBeth case leads to the rather implausible finding that all the 27 test portfolios display large and negative average abnormal returns that cannot be justified by exposure to systematic risks. This means that all of our portfolios would have been systematically and persistently over-priced during our sample period. This would be an overwhelming indication of irrational exuberance of the U.S. market over a 31-year long sample. On the contrary, in the Bayesian case, the values of the posterior means for the parameters capturing any mispricings as well as their signs are economically plausible.¹ Interestingly, this result and the one above, are attained by a model that turns out to be more parsimonious than the rolling window, two-stage Fama–MacBeth (one may say, seminonparametric) implementation because our latent mixture framework restricts both risk exposures and volatilities to change under a change-point structure, while in the two-pass approach, even though the change is limited, coefficients change in each time period for sure.

Third, under a Fama–MacBeth implementation, few of the factor unit risk premia are precisely estimated while a few carry the wrong sign. The only quantity for which there is compelling evidence is the parameter that measures average cross-sectional abnormal returns not justified by risk exposures. In a sense, all that a standard Fama–MacBeth approach reveals is that U.S. asset data contain strong evidence of structural mispricings. On the opposite the Bayesian estimates of the risk premia are considerably more stable and, more importantly, a few of them are precisely estimated. Here one result is striking: with reference to the full sample, market and real consumption growth risks produce the only precisely estimated risk premia and display the expected positive sign. Also, the Bayesian design gives evidence of moderate but precisely estimated mispricings of 0.37% per month. However such mispricing at least fail to consistently appear in sub-sample analyses.

Fourth, we compute predictive log-likelihood (PL-L) scores to compare the four alternative implementations entertained in this paper to find that the model yielding the highest PL-L is the full model with time-varying factor exposures and stochastic volatility both subject to discrete breakpoints. The full model outperforms not only the Fama–MacBeth two-step approach but also restricted version of the latent mixture, such as the homoskedastic and continuous breakpoint cases. There is also clear structure in the ability of BTVSVB model to outperform the Fama–MacBeth implementation: this occurs with remarkable strength for the largest cap decile equity portfolios (8–10) and for Treasuries (both intermediate and long term), where the improvement in the recorded PL-L exceeds 10%.

Although the issue of searching the most appropriate methods to learn about the pricing kernel that underlies the observed cross-section of asset returns using mixture models with latent breaks are new to the best of our knowledge, Ouyssse and Kohn (2010)

¹ Even more interestingly, once we take into account their 90% confidence bands, it is clear that in 27 out of 27 portfolios, the reported posterior means fail to be precisely estimated, in the sense that the 90% band always includes a zero abnormal return during the period 1980–2010.

متن کامل مقاله

دریافت فوری ←

ISIArticles

مرجع مقالات تخصصی ایران

- ✓ امکان دانلود نسخه تمام متن مقالات انگلیسی
- ✓ امکان دانلود نسخه ترجمه شده مقالات
- ✓ پذیرش سفارش ترجمه تخصصی
- ✓ امکان جستجو در آرشیو جامعی از صدها موضوع و هزاران مقاله
- ✓ امکان دانلود رایگان ۲ صفحه اول هر مقاله
- ✓ امکان پرداخت اینترنتی با کلیه کارت های عضو شتاب
- ✓ دانلود فوری مقاله پس از پرداخت آنلاین
- ✓ پشتیبانی کامل خرید با بهره مندی از سیستم هوشمند رهگیری سفارشات