Journal of Empirical Finance 16 (2009) 537–556

Contents lists available at ScienceDirect

Journal of Empirical Finance

journal homepage: www.elsevier.com/locate/jempfin

# Learning about beta: Time-varying factor loadings, expected returns, and the conditional CAPM[☆]

Tobias Adrian [a,1], Francesco Franzoni [b,*]

[a] *Federal Reserve Bank of New York, 33 Liberty Street, New York, NY 10045, United States*
[b] *Swiss Finance Institute, University of Lugano, Via Buffi 13, Lugano, 6904, Switzerland*

ARTICLE INFO

*Article history:*
Received 28 November 2006
Received in revised form 3 February 2009
Accepted 17 February 2009
Available online 9 March 2009

*JEL classification:*
G12
C11

*Keywords:*
Beta
CAPM
Kalman filter
Anomalies
Value premium

ABSTRACT

We amend the conditional CAPM to allow for unobservable long-run changes in risk factor loadings. In this environment, investors rationally "learn" the long-run level of factor loadings from the observation of realized returns. As a consequence of this assumption, we model conditional betas using the Kalman filter. Because of its focus on low-frequency variation in betas, our approach circumvents recent criticisms of the conditional CAPM. When tested on portfolios sorted by size and book-to-market, our learning-augmented conditional CAPM passes the specification tests.



Fama and French (1992) present compelling evidence that the unconditional capital asset pricing model (CAPM) does not account for returns on portfolios sorted by size and book-to-market (B/M). Since then, the literature on asset pricing has developed alternative theories, which depart from the original model along several dimensions. Promising avenues of research, which preserve the single-factor structure, have been conditional versions of the CAPM. The idea behind this approach is that, even though CAPM holds conditionally on time $t$ information, it does not hold unconditionally. Accordingly, the poor empirical performance of CAPM might be due to its failure to account for time variation in conditional moments.

Among the many studies of the conditional CAPM, recent implementations are put forward by Jagannathan and Wang (1996), Ferson and Harvey (1999), and Lettau and Ludvigson (2001). Earlier papers include Ferson, Kandel, and Stambaugh (1987), Bollerslev, Engle, and Wooldridge (1988), and Harvey (1989). In all cases, the authors explicitly model the evolution of the conditional distribution of returns as a function of lagged state variables. They specify the covariance between the market return and portfolio returns as affine functions of these variables. This specification is estimated as a multifactor model, in which the additional factors are the interactions between the market return and the state variables.

A recent paper by Lewellen and Nagel (2006) casts doubts on the empirical success of this approach. While acknowledging that betas vary considerably over time, these authors suggest that the covariation between sample estimates of betas and the market risk premium is not large enough to justify the deviations from the unconditional CAPM observed for value and momentum

☆ We are grateful to Manuel Amador, Olivier Blanchard, Xavier Gabaix, Jonathan Lewellen, Sendhil Mullainathan, and Steve Ross for comments. We would like to thank the associate editor Zhenyu Wang and an anonymous referee for important contributions. The views expressed in this paper are those of the authors and do not necessarily reflect the position of the Federal Reserve Bank of New York or the Federal Reserve System.
* Corresponding author. Tel.: +41 58 666 4117; fax: +41 58 666 4734.
*E-mail addresses:* tobias.adrian@ny.frb.org (T. Adrian), francesco.franzoni@unisi.ch (F. Franzoni).
[1] Tel.: +1 212 720 1717; fax: +1 212 720 1582.

0927-5398/$ – see front matter © 2009 Elsevier B.V. All rights reserved.
doi:10.1016/j.jempfin.2009.02.003

portfolios (Fama and French, 1993; Jegadeesh and Titman, 1993). They argue that the good empirical performance of previous conditional studies is due to their cross-sectional design—which ignores key theoretical restrictions on the estimated slope coefficients—and suggest time-series regressions instead.

In this paper, we complement the conditional CAPM literature by modeling a new type of time variation in conditional betas. There is substantial evidence that the risk of some asset classes has experienced long-run movements. For example, CAPM regressions on CRSP data suggest that value stocks had higher betas than growth stocks between 1926 and 1963. The opposite is true in the later sample.[2] This low-frequency evolution is different from the covariance of beta with the equity premium that is typically modeled in conditional CAPM applications.

Our goal is to explore the implications of long-run changes in factor loadings for the tests of conditional models. We expect this effect to be important for the assets that have displayed major swings in their betas, such as portfolios sorted by size and B/M.

We tackle this issue by assuming that betas change over time following a mean-reverting process. Also important, we postulate that investors are unaware of the long-run level of a given asset's risk. Hence, they need to infer both the level of beta and its long-run mean from the history of realized returns. Through a stylized model, in which the conditional CAPM holds, we make the point that the risk loading that determines expected returns is the expectation of beta that results from the learning process. We model this expectation via a Kalman filter, in which the factor loading is treated as a latent variable, and argue that betas should be estimated accordingly in empirical applications. Consistent with this theoretical argument, we use Kalman-filtered betas to explain the returns of the twenty-five portfolios sorted by size and B/M.

Our first set of results concerns the CAPM, augmented with learning, but without conditioning variables. When we use a mispricing measure borrowed from Campbell and Vuolteenaho (2004), the sole contribution of learning is a reduction of about 45% in mispricing relative to the unconditional CAPM. Low-frequency movements of beta play a crucial role in this result. Investors' inferences about the long-run level of beta can cause a significant difference between the ex-ante expected level of risk and ex-post estimates from typical OLS regressions. This mechanism is particularly relevant for portfolios of value and small stocks that have experienced considerable long-run variation in beta. We argue that the wedge between investors' ex-ante expectations of beta and ex-post OLS estimates can account for a large fraction of the unconditional alpha in standard OLS time-series regressions. In other words, the mismeasurement of expectations of beta and, hence, of equilibrium expected returns can be the source of the apparent mispricing.

Our second set of results concerns the conditional CAPM with different sets of scaling variables. We confirm Lewellen and Nagel's (2006) finding that conditioning variables alone do not improve the performance of CAPM much in time-series tests. However, once we introduce learning in conditional CAPMs, the model is no longer rejected, and the aggregate pricing error is reduced by 45% (it decreases by 65% when compared to the unconditional CAPM). In every specification, we find that the model without learning is rejected, whereas the model with learning is not. Pricing errors of our (one-factor) learning-augmented conditional CAPM are comparable in magnitude to the ones from the Fama–French three-factor model.

The intuition behind the empirical success of our learning-augmented CAPM in pricing B/M- and size-sorted portfolios is as follows. Consider the value premium. The failure of CAPM to price value stocks in an OLS framework is due to value portfolios having high average returns but low estimated betas. In our approach based on the Kalman filter, given the decrease in systematic risk of value stocks, the high level of the past factor loading affects today's estimates and makes them larger than OLS betas. A high estimate of beta is thus matched with high average returns, and the estimated alpha of value stocks is reduced. A similar intuition applies to small stocks, which have also experienced a decline in systematic risk. The situation is reversed in the case of large and growth stocks, whose levels of risk have increased over time.

Our results provide a support for the conditional CAPM that is not subject to Lewellen and Nagel's (2006) critique. The type of variability in beta that affects our estimates is low-frequency, idiosyncratic variation, experienced over a long time horizon. Lewellen and Nagel, instead, focus on the cyclical covariation between beta and market risk premium. Consistent with their prediction, we find that the standard conditioning variables do not lead to success of CAPM when the model is tested in the time series via OLS. However, the inclusion of these conditioning variables in our learning framework reestablishes the success of the conditional CAPM in pricing size- and B/M-sorted portfolios.

Following much of the literature on asset pricing since the early 1970s, our test assets are portfolios of stocks. In the context of conditional asset pricing models, this procedure implicitly assumes that investors price stocks according to the risk of the asset class to which they belong. The asset class is defined by the relevant characteristic according to which portfolios are sorted. In our empirical application, we use the Kalman filter to estimate betas of size- and B/M-sorted portfolios rather than individual stocks. In this sense, we proceed as if investors estimated a unique beta for all the stocks in the same portfolio. While this assumption has the flavor of some "behavioral" specification, such as "style investing" (Barberis and Shleifer, 2003), it can also be motivated on rational grounds and anecdotal evidence. First, it makes sense to believe *a priori* that when the time series of a firm's returns is not long enough to allow a reliable estimation of its beta, the rational way to predict its risk is to compare the company to other firms with similar characteristics, for which a longer time series is available. Second, we believe the spirit of this assumption is in line with an important feature of the actual investment process. For example, Barra, a provider of beta estimates, reports a "fundamental measure" of a stock's beta. Its estimate is the weighted average of the betas of a set of characteristic-based portfolios to which a stock belongs.[3] Our empirical focus on the size and B/M characteristics is a clear simplification of the multidimensionality of the problem. Still, we feel that we replicate a salient feature of institutional investors' approach in computing betas.

[2] See, for example, Franzoni (2002, 2009), Campbell and Vuolteenaho (2004), Fama and French (2006), and Ang and Chen (2007).

[3] Barra says that this fundamental beta is superior to the historical beta in predicting future risk. See, for example, Barra's newsletter from September 1988: "What's New About Beta?"