How Do Alphas and Betas Move?
Uncertainty, Learning and Time Variation in Risk Loadings

CARMINE TRECROCI*
Department of Economics, University of Brescia

December 2009

Abstract

I employ a parsimonious model with learning but without conditioning information to extract time-varying measures of market-risk sensitivities, pricing errors and pricing uncertainty. Parameters estimated for U.S. equity portfolios show significant fluctuations, along patterns that change across size and book-to-market categories of stocks. Time-varying betas display superior predictive accuracy for portfolio returns against constant and rolling-window OLS estimates. I also study the relationship of betas with business-cycle variables, finding that those of high BE/ME stocks move pro-cyclically, unlike those of low BE/ME stocks. Investment growth, rather than consumption, predicts the betas of high BE/ME and small-firm portfolios.

Keywords: Conditional CAPM, Time-varying beta, Uncertainty, Learning.
JEL Codes: G12, C51.

1 Introduction

According to the Capital Asset Pricing Model (CAPM), differences between asset expected returns reflect differences in their exposure to systematic risk, as measured by market betas. In practice, market-risk sensitivities are the slope coefficients from OLS regressions of asset excess returns on the market's excess return. However, the descriptive accuracy of this procedure rests on the key assumption that the actual parameters are time-invariant. Moreover,

*Department of Economics, University of Brescia,Via San Faustino 74/B, 25122 Brescia. Tel. +39 030 2988812, fax +39 030 2988837, trecroci@eco.unibs.it, www.eco.unibs.it/~trecroci. This paper was in part developed while I was a visiting scholar at the Department of Finance of the Carroll School of Management, Boston College. I wish to thank Pierluigi Balduzzi for very useful suggestions on this research. All errors are mine.

there are some reasons, both theoretical and practical, which instead call for a conditional specification of the CAPM. In particular, one could allow for parameter uncertainty and variation over time. In this paper, I construct and estimate a parsimonious one-factor model with time-varying alphas and betas that are endogenous with respect to the uncertainty surrounding their actual values. The market-risk sensitivities that I obtain reveal superior predictive ability for portfolio returns against constant and rolling-window OLS estimates. I also evaluate their fluctuations, finding that alphas and betas of portfolios characterised by different book-to-market values (BE/ME) and market capitalization (size) evolve according to different cyclical patterns.

There are no reasons to think that the exposures to fundamental sources of risk vary across assets, i.e. across claims to different cash flows, but not over time, i.e., when the information set and economic circumstances affecting the valuation of the firm's cash flows possibly change. In addition, if the variance of market return or its covariance with asset returns is time-varying, an asset's beta with the market will change over time too. Indeed, there is ample evidence on the persistence and heteroskedasticity of market returns at business-cycle frequencies (Schwert, 1989a, b). Finally, even assuming that individual stocks have time-invariant betas, changes in portfolio weights imply that portfolio returns satisfy a linear factor model, but one with time-varying coefficients and a heteroskedastic disturbance term (see Mamaysky et al., 2008).

Many real-world factors are likely to play a significant role in the determination of market betas. One of them is investors' uncertainty, which might affect investment choices in various ways. This study focuses on broad economic uncertainty, that is, on the impact that volatility and heteroskedasticity of fundamentals have on investors' ability to identify the distribution of asset payoffs. It is plausible to think that, under uncertainty, investors' forecasts of key quantities like market betas will be the result of some complex learning process that reflects that uncertainty. Indeed, puzzles and anomalies are pervasive in empirical asset pricing. At least in part this evidence might be due to the fact that some key parameters of financial models are in fact uncertain and subject to learning effects. For instance, Bekaert et al. (2009) study the relative importance of variations in heteroskedasticity of fundamentals and stochastic risk aversion on various asset prices and returns[1]. Pástor and Veronesi (2009) argue that the market betas of innovative firms are likely to increase during technological revolutions. These and other arguments suggest that the traditional, constant-coefficient models for portfolio returns could mis-specify the identification of asset risk. Changes in the structure of the economy and in financial markets make reasonable to model risk sensitivities as potentially time-varying quantities, particularly over long samples and at lower frequencies.

In this paper, I propose and estimate a conditional one-factor model with time-varying alphas and betas, and extract an endogenous measure of the uncertainty surrounding their values. In detail, uncertainty is defined as the conditional error variance of the optimal forecast of alphas and betas. The aim is to account for the effects of uncertainty and change by replicating the learning process of rational investors. I assume that the latter must infer the risk loadings from available information, and optimally update them as new information becomes available. This paper also holds that changes in market returns effectively summarize the arrival of relevant information; hence, the estimated risk loadings do not rely on conditioning information.

Accordingly, I design a parsimonious model that allows for changes in perceived risks due to factors fully unobserved by the econometrician, such as shifts in the quantity of market risk that might be learning-induced. I then estimate a conditional one-factor relationship with time-varying parameters, based on the Kalman filter. This methodology generates monthly alpha and beta time series without relying on conditioning information or time/frequency assumptions. The main advantages of this estimation strategy over existing alternatives are its simplicity, and its ability to adapt to assets' or portfolios' actual loadings on market risk in a way that constant-coefficient, but also rolling or fixed-window OLS regressions, simply do not permit to. In contrast to some recent contributions, the resulting joint estimates of each

[1] They also find a limited but positive effect of shocks to the conditional volatility of the dividend growth process on risk premia (which are mainly driven by shocks to risk aversion).

period's conditional alphas and betas are obtained without making any assumption about period-to-period variation in beta. Consequently, the time-varying betas in this paper denote systematically superior predictive ability for portfolio returns against conventional rolling-window OLS estimates. Furthermore, estimated parameters fluctuate significantly over time; this confirms that investors update their forecasts on a more frequent and systematic basis than existing analyses entertain. Finally, I study whether market-risk sensitivities evolve according to some cyclical pattern, finding clear-cut evidence that their relationship with the business cycle depends persistently on portfolio characteristics such as size and book-to-market.

Methodologically, the determination of CAPM coefficients should be endogenous with respect to investors' uncertainty about actual factor loadings. This means that the variance of returns should be time-varying too. The empirical exercise in this paper accounts for two sources of uncertainty: uncertainty arising from future idiosyncratic risk, and uncertainty arising because of evolution in the risk loadings. Conditional uncertainty is therefore directly associated to observed returns, which contain and update the information relevant for investment choices. In addition, the model allows for both time variation in the mean and homoskedastic stochastic components of the alpha and beta processes, thus combining features that the existing literature does not consider jointly[2].

I apply this methodology to the 1926-2007 monthly returns of U.S. equity portfolios sorted by size and BE/ME, obtaining estimates of alphas, betas and pricing uncertainty that evolve over time. Crucially, the time-varying, Kalman-filter based (TVK, henceforth) estimates that I obtain depend only on portfolio and market returns and appear to be precisely estimated. To assess their predictive accuracy against conventional rolling OLS betas, I perform an out-of-sample simulation over a hedging strategy based alternatively on these two measures.

A large literature explains the risk premia associated with size and BE/ME ratios with the relationship between those characteristics and fluctuations in aggregate consumption or wealth. I further examine whether the TVK alphas and the risk loadings of portfolios that

[2] Jostova and Philipov (2005) is the only exception.

belong to different BE/ME and size categories evolve according to different patterns. Understanding exactly how the sensitivity of characteristic portfolios to market risk changes over time, and in relation to which economic conditions, is an issue of primary interest. Accordingly, I evaluate the association of TVK betas with key state variables and macroeconomic indicators, obtaining fresh evidence on the dynamics of market-risk sensitivities at business-cycle frequency.

The remainder of the paper and its main findings are as follows. Next Section sets out the literature background to the empirical exercise in this paper. In Section 3 I test formally for the presence of breaks in time-invariant CAPM coefficients, detecting multiple structural breaks. Section 4 introduces a specification of the Kalman filter that accounts for the learning problem of investors under uncertainty. Section 5 presents estimates of time-varying alphas, betas and pricing uncertainty, and tests for the relative information content of TVK betas for market returns. Market betas are more tightly estimated and have superior predictive ability for actual portfolio returns than those obtained through the conventional rolling-window approach. Also, the evolution of TVK parameters appears to be very rich and differentiated across portfolios sorted on the basis of book-to-market ratios and market capitalization. Section 6 evaluates the association of time-varying betas with business-cycle indicators at the monthly and quarterly frequencies. The betas of high BE/ME stocks turn out to move pro-cyclically, whereas those of low BE/ME stocks denote an opposite, though weaker tendency. Large-cap portfolios have betas strongly correlated with state variables, and which help predict future output. Finally, investment growth, rather than consumption growth, helps forecast the betas of high BE/ME and small firms portfolios, whereas risk loadings of large-cap portfolios anticipate future output developments. These results lend some support to recent production-based asset-pricing models. Section 7 concludes. The Appendix contains summary statistics about the returns and CAPM parameters of the test portfolios.

2 Literature review

There are essentially two sources of systematic risk. First, the risk of an asset is a function of the sensitivity of its cash flows to fluctuations in the market return or to changes in the economy's rate of growth. Also, the present value of those cash flows is contingent on the aggregate discount rate. Depending on the time distribution of cash flows, shocks to the discount rate drive changes in asset returns. Hence, the sensitivity to cash-flow risk and discount-rate risk determines an asset's risk-return trade-off[3].

There is substantial evidence on the variability over time of market premia (see for instance, Ang and Bekaert, 2007; Cooper and Priestley, 2009), and some also of similar behaviour by market betas (Lewellen and Nagel, 2006; Ang and Chen, 2007; Adrian and Franzoni, 2009). Therefore, the well-known poor empirical performance of unconditional CAPM might be the result of time variation in the conditional moments used to capture systematic risk (see Lettau and Ludvigson, 2001; Santos and Veronesi, 2004; Zhang, 2005; Lewellen and Nagel, 2006). A common approach to testing conditional CAPM is to model betas as a function of observed macroeconomic and financial variables. However, these tests are strictly valid only if the econometrician knows the full set of state variables available to investors. Lewellen and Nagel (2006) use short-window regressions to estimate conditional alphas and betas for momentum, size and BE/ME portfolios. They find that the measured variation in CAPM coefficients is not sizeable enough to explain large unconditional alphas. Lewellen and Nagel also uncover large conditional pricing errors, hence validating once again the model's poor empirical performance. Fama and French (2006) employ rolling OLS regressions to estimate one-year betas and alphas, whose properties essentially confirm the existence of value, size and momentum "anomalies"[4]. On the other hand, Ang and Chen (2007) extract conditional CAPM parameters using a model with persistent betas, time-varying risk premia and sto-

[3] Campbell and Vuolteenaho (2004) and Bansal et al. (2005) among others study the relevance of shocks to cash flows and the discount rate in determining the cross-section of returns. Santos and Veronesi (2004) build a general-equilibrium model that explains the time variation of betas.

[4] Fama and French (2006) acknowledge that betas should not be modeled as time-invariant; they use slope dummies to model discrete breaks in betas of coarser portfolio sorts than those I use here.

chastic systematic volatility. Their results support the CAPM over a long sample span, and therefore clash with the aforementioned evidence.

The empirical exercise in this paper, which does not focus on a formal test of CAPM, shares the tenet of the literature that latent measures of systematic risk are likely to vary significantly over time. However, the work in this article goes beyond the existing literature along various dimensions.

First, the one-factor model that I estimate yields parameters that account endogenously for the level of uncertainty. In particular, this paper defines conditional uncertainty as the variance of the error in the optimal forecast of alphas and betas. Pástor and Veronesi (2009) argue that investors' uncertainty about future cash flows warrants a larger role for learning mechanisms in various areas of finance, starting with stock valuation. Besides the paper by Bekaert et al. (2009), other recent studies take this view. For instance, Ozoguz (2009) extracts some measures of uncertainty based on two-state regime-switching models for market return and aggregate output. She finds that uncertainty is a priced risk factor in the cross-section of stock returns[5]. This paper has a different and complementary aim, and adopts a specification of the market model that allows for multiple sources of uncertainty about market risk. In addition, my approach models investor's learning about unobserved conditional moments in a parsimonious way, yields minimum mean-squared-error forecasts of the quantity of risk, and it is based only on observed returns.

A second improvement in relation to existing studies derives from the use of the Kalman filter. The algorithm provides consistent parameter estimates and therefore allows for robust inference. On the contrary, the inference generally based on betas estimated through short-window or rolling regressions rests on a key hypothesis. The variance of the rolling beta is generally held (for instance in Fama and French, 1997) to be the sum of the variance of the true market beta and the variance of the estimation error. But this holds true only if the sampling error of the market beta is uncorrelated with the true value of the beta. Now, if investors formulate their forecasts under uncertainty, the volatility of beta's estimation error

[5] See Anderson et al. (2009) for an interesting alternative approach.

is likely to be increasing in the volatility of the true beta, making such assumption quite restrictive. Also, holding that coefficients are constant in-sample but changing discretely out-of-sample introduces a small-sample bias that cannot be corrected using HAC estimators. Lewellen and Nagel (2006) and Campbell and Vuolteenaho (2004) employ short-window regressions. For instance, the former employs daily returns to estimate monthly betas. This could generate some bias in the estimates, as a consequence of the assumption about period-to-period variation in beta and various microstructure issues (bid/ask bounce, irregular trading, stale pricing). Bali (2008) employs a GARCH (1,1) model to extract conditional moments and betas for a variety of test portfolios, but using conditioning information. On the contrary, Jostova and Philipov (2005) model time-varying betas of industry portfolios using Bayesian techniques (Monte Carlo Markov Chain with Gibbs sampling) based on an explicitly mean-reverting beta. Their study provides time-varying betas that are shown to be more precisely estimated than those from either rolling regressions or GARCH approaches[6]. In this paper, I directly estimate model parameters that share the returns' frequency, without using state variables. An approach that avoids conditioning information has the advantage of deriving betas directly from portfolio returns, thereby reducing any potential omitted-variable bias due to mis-specification of the information set.

Adrian and Franzoni (2009) too employ a learning-based form of CAPM and estimate betas using the Kalman filter (albeit on quarterly data). However, they assume constant long-term means for asset betas, estimates are based on conditioning variables, and their model implies that, asymptotically, investors' uncertainty over the value of beta is solved, i.e., it disappears. These hypotheses are not consistent with the idea that equity risk changes over time, and that investors face uncertainty on a systematic basis. Still, Adrian and Franzoni's alphas and idiosyncratic risk are not time-varying, whilst it would be appropriate to endogenize their dynamics along that of betas. Finally, their use of the smoothing version of the filter likely makes estimates further suffer from a look-ahead bias. Indeed, an additional

[6] However, conditional alphas and idiosyncratic risk are not endogenously determined in Jostova and Philipov (2005), and their estimates show that a mean-reversion assumption about betas might significantly constrain the dynamics of factor loadings.

difference over previous contributions is that in this empirical exercise I estimate jointly time-varying betas, alphas and a measure of pricing uncertainty.

3 Are market factor loadings time-invariant?

As reported above, Jostova and Philipov (2005) and Ang and Chen (2007) find substantial evidence of inconsistence in the estimation of conditional CAPM coefficients from constant-parameter and rolling regressions. Building on such results, in this Section I check for the presence of breaks in the parameters of conventional one-factor models, with the help of formal stability tests.

In this paper, the test assets are the well-known 5×5 Fama-French (FF) value-weighted portfolios, formed as the intersections of independent sorts of stocks on the basis of book-to-market ratios and market capitalization, spanning from July 1926 to August 2007[7]. There are at least two motivations for the use of these returns. First, starting from Fama and French (1992), these portfolios have featured in several studies closely related to the objective of this paper. Second, estimated betas are inevitably measured with error, and this error typically induces a downward bias in regression betas. Hence, using portfolio returns rather than those of individual stocks should help reduce this bias. Portfolio betas are likely to be more precisely measured also because portfolio returns have lower idiosyncratic risk.

Let us begin with a look at the data. The Appendix at the end of the paper reports some descriptive statistics for the average returns on the test assets, and estimates from the

[7] Data are from the CRSP database. The portfolios, which are constructed at the end of each June, are the intersections of 5 portfolios formed on size (market equity, ME) and 5 portfolios formed on the ratio of book equity to market equity (BE/ME). The size breakpoints for year t are the NYSE market equity quintiles at the end of June of t. BE/ME for June of year t is the book equity for the last fiscal year end in $t-1$ divided by ME for December of $t-1$. The BE/ME breakpoints are NYSE quintiles. The portfolios for July of year t to June of $t+1$ include all NYSE, AMEX, and NASDAQ stocks for which there are market equity data for December of $t-1$ and June of t, and (positive) book equity data for $t-1$. The ending of the sample right before the start of the 2007-2009 crisis is fully unintentional. Nevertheless, this allows looking for possible signs of the impending financial turmoil and calls for further investigation based on data covering its entire duration.

I thank Kenneth French for making a large amount of data publicly available in his online data archives, from which I downloaded the data I use in this paper.

standard one-factor regression:

$$R_t^{ei} = a^i + b^i R_t^{eM} + e_t^i,$$

where $R_t^{ei} = R_t^i - R_t^f$ is the return on test portfolio i in excess of the one-month Treasury bill rate and $R_t^{eM} = R_t^M - R_t^f$ is the excess return on the market, i.e., the value-weighted return on all NYSE, AMEX, and NASDAQ stocks minus the one-month Treasury bill rate (from Ibbotson Associates). Each portfolio is dubbed by reference to its BE/ME and size category. In the notation, firm size increases from S1 to S5 and the BE/ME ratio from B1 to B5[8].

The data (see Table A1) confirm the existence of a significant value premium in both samples, and of a size premium particularly in the longest one. The focus of this paper is not on whether the value premium is definitely limited to small-stock returns, but we note that average returns within each BE/ME class are smaller for stocks of larger firms, and for low BE/ME portfolios within each size category. As in other studies, portfolio B1S5 breaks such regularities, especially over the early part of the sample. In fact, this portfolio has even missing observations from July 1930 to June 1931[9].

Over the whole sample, most regression intercepts are positive and sometimes statistically significant, particularly for high BE/ME portfolios[10]. Overall, regression slopes appear to be very tightly estimated. That said, Figure A1 gives some perspective on the extent to which unconditional market betas fail to explain the so-called value and size premia. Average returns are not only too loosely associated with market betas; the lines connecting portfolio

[8] So, for instance, portfolio B2S5 contains stocks of the largest firms within the second highest BE/ME category, and so on. In addition to the 25 size-BE/ME portfolios, I report results for some of their combinations. $H - L$ is the difference between the average return for the two highest BE/ME quintiles within a size quintile and the average return of the two lowest BE/ME quintiles. Similarly, $S - B$ is the difference between the average return for the two smallest quintiles within a BE/ME quintile and the average return of the two largest quintiles within a BE/ME quintile.

[9] I estimate the regressions over both the whole sample (1926-2007) and the shorter period 1963-2007. Tables A2 and A3 exhibit estimated coefficients a and b and their t-values, constructed using HAC standard errors, following Andrews (1991).

[10] For the 1963-2007 sample, one obtains broadly similar estimates, with somewhat more significant and sizeable high BE/ME intercepts, but also a few negative and statistically significant ones for the low BE/ME portfolios.

returns in each panel also show that differences in BE/ME and size yield to differences in average returns that are negatively and positively related, respectively, to regression slopes. This confirms the existence of non-beta risks related to value and size that go unaccounted for by the way in which the unconditional one-factor model describes average returns across categories of stocks[11].

Next, I investigate whether market-risk sensitivities are stable throughout the sample. I employ the test of multiple structural changes devised by Bai and Perron (1998, 2003; see also Qu and Perron, 2006). The test is particularly suited to our case, as it is designed for the evaluation of multiple structural breaks in the context of linear models estimated by OLS. It provides a way of testing, through the use of a Sup Wald-type test, the null hypothesis of no change in the coefficients versus an alternative containing a finite number of shifts. Furthermore, Bai and Perron propose an algorithm that allows testing the null of m changes versus the alternative hypothesis of $m+1$ changes. I apply the tests to a specification with a constant, also accounting for serial correlation and different variances in the residuals[12]. For brevity, Table I shows the results for five "average" portfolios formed on the ratio of book equity to market equity and five "average" portfolios formed on size, over the period 1963-2007[13].

Only for few cases the $\sup \mathbf{F}_T(m)$ statistics, computed by pre-specifying the number of breaks, are not significant. Also, for the $\mathbf{UD}\max$ and $\mathbf{WD}\max$ tests, which allow testing the null of no structural break against an unknown number of breaks, most statistics are significant at conventional levels, thus revealing multiple structural breaks in the coefficients of the one-factor regression. Therefore, traditional OLS regressions of the one-factor model that are based on time-invariant coefficients are definitely mis-specified and produce inconsistent inference about the distribution of risk premia. Next Section introduces an intuitive framework that models risk loadings as time-varying processes.

[11] Fama and French (2006) elaborate more extensively on this point.

[12] I thank Pierre Perron for making available his GAUSS procedure to implement the test.

[13] The results of the test for the longer 1926-2007 sample and all 25 FF portfolios are in line with those reported here.

4 Uncertainty, learning and time variation in a parsimonious model

The arrival of new information and the presence of investor uncertainty have defined implications on CAPM predictions. The conditional CAPM states that asset expected returns are proportional to their conditional market betas (Zhang, 2005). I define the latter under uncertainty as the expectation of the factor loadings on the market:

$$\mathbb{E}\left[R_t^{ei}\right] = \mathbb{E}\left[R_t^{eM}|\psi_{t-1}\right] \cdot \mathbb{E}\left[\beta_t^i|\psi_{t-1}\right] \tag{1}$$

Here R_t^{ei} is asset i's excess return, β_t^i is CAPM's beta and ψ is the information set. Accordingly, I posit that returns follow a linear regression model, in which the coefficients change over time according to an autoregressive dynamics:

$$R_t^{ei} = x_t\beta_t^i + \varepsilon_t^i, \quad t = 1, 2, \ldots, T \tag{2}$$

$$\beta_t^i = \widetilde{\beta}^i + F^i\beta_{t-1}^i + \upsilon_t^i \tag{3}$$

where

$$\varepsilon_t^i \sim IIDN\left(0, S\right) \tag{4}$$

$$\upsilon_t^i \sim IIDN\left(0, Q\right) \tag{5}$$

and x_t is a vector of exogenous or predetermined variables. Importantly, ε_t^i and υ_t^i are mutually independent and x_t contains only a constant and the market's excess return. Unlike in Adrian and Franzoni's (2009) setup, the betas (and alphas) are not held to be conditional on any exogenous variable. To spare notation, I drop the superscript "^{i}" to denote asset i's return, CAPM coefficients, etc.

If investors were fully informed and under no uncertainty, all parameters $(\widetilde{\beta}, F, S, Q)$ would be known. If this was really the case, a sequence of GLS regressions would deliver an estimate of the state vector. However, such approach tends to be extremely inefficient in terms of computational burden. More important, we have discussed above several reasons for which it is plausible to think that investors are instead uncertain about the true values of those parameters and therefore need to update systematically their forecasts. Finally, if only some of the hyperparameters were not known, they would have to be estimated anyway before making any inference about β_t. All this leads quite naturally to consider the Kalman filter (KF, henceforth) to make inferences about β_t.

The Kalman filter is a recursive procedure for computing the estimator of a time-t unobservable component, the state vector, based only on information available up to time t. When the shocks to the model and the initial unobserved variables are normally distributed, the KF also enables to compute the likelihood function through prediction error decomposition[14]. Both features are particularly suited for the treatment of time variation and uncertainty in the cross-section of market risk.

The KF computes a minimum mean-squared-error estimate of β_t conditional on ψ. Depending on the information set used, one obtains filtered or smoothed estimates. The filter, which is used in this paper, refers to an estimate of β_t based on information available up to time t, whereas the smoothing version of the Kalman algorithm yields an estimate of β_t based on all the available information in the sample through time T. The latter version, which is employed in Adrian and Franzoni (2009), hence assumes that investors know the true value of hyperparameters -like the long-run level of beta- when they form forecasts of time-varying parameters.

Now, I define $\beta_{t|t-1} = \mathbb{E}\left[\beta_t|\psi_{t-1}\right]$ as the expected value of β_t conditional on ψ_{t-1}, whereas $\beta_{t|t} = \mathbb{E}\left[\beta_t|\psi_t\right]$ represents the estimate of β_t conditional on ψ_t, and therefore on the realization of the prediction error. Let us also define $P_{t|t-1} = \mathbb{E}\left[\left(\beta_t - \beta_{t|t-1}\right)\left(\beta_t - \beta_{t|t-1}\right)'\right]$, and $P_{t|t} = \mathbb{E}\left[\left(\beta_t - \beta_{t|t}\right)\left(\beta_t - \beta_{t|t}\right)'\right]$ as their respective covariance matrices[15].

The asset's expected excess return, which represents an optimal forecast given information up to time $t-1$, is $R^e_{t|t-1} = \mathbb{E}\left[R^e_t|\psi_{t-1}\right] = x_t\beta_{t|t-1}$. This forecast has prediction error equal to $\eta_{t|t-1} = R^e_t - R^e_{t|t-1}$, in turn characterised by conditional variance $f_{t|t-1} = \mathbb{E}\left[\eta^2_{t|t-1}\right]$.

[14] See Hamilton (1989) and Kim and Nelson (1999) for reviews of the Kalman filter.

[15] For smoothed Kalman estimates, we have $\beta_{t|T} = \mathbb{E}\left[\beta_t|\psi_T\right]$ and $P_{t|T} = \mathbb{E}\left[\left(\beta_t - \beta_{t|T}\right)\left(\beta_t - \beta_{t|T}\right)'\right]$, respectively.

The timing assumption of the model is straightforward. At the beginning of time t, x_t becomes available; at the end of time t a new realization of R_t^e becomes public knowledge. As a consequence, the basic KF involves two steps:

Step 1 At the beginning of time t, investors formulate an optimal prediction of the asset's expected return, $R_{t|t-1}^e$, based on all the information up to time $t-1$. To do this, investors need to compute $\beta_{t|t-1}$:

$$\beta_{t|t-1} = \widetilde{\beta} + F\beta_{t-1|t-1} \tag{6}$$

$$P_{t|t-1} = FP_{t-1|t-1}F' + Q \tag{7}$$

$$\eta_{t|t-1} = R_t^e - R_{t|t-1}^e = x_t\beta_{t|t-1} \tag{8}$$

$$f_{t|t-1} = x_tP_{t|t-1}x_t' + S \tag{9}$$

Besides $\beta_{t|t-1}$, the KF algorithm therefore generates an estimate of the conditional variance of the forecast errors, eq. (9). This equation demonstrates that the model accounts for two sources of uncertainty: uncertainty arising from future idiosyncratic risk, and uncertainty arising because of evolution in the model's coefficients. In this simple way, conditional uncertainty is directly associated with observed returns, which contain and update the information relevant for investment choices.

Step 2 Once R_t^e is realized at the end of time t, the prediction error can be calculated: $\eta_{t|t-1} = R_t^e - R_{t|t-1}^e$. This contains new information about the model's coefficients β_t, beyond that contained in $\beta_{t|t-1}$. Thus, after observing R_t^e, a more accurate inference about β_t can be made. $\beta_{t|t}$, an inference of β_t based on information up to time t, has the following form

$$\beta_{t|t} = \beta_{t|t-1} + K_t\eta_{t|t-1} \tag{10}$$

$$P_{t|t} = P_{t|t-1} - K_tx_tP_{t|t-1} \tag{11}$$

The quantity

$$K_t = P_{t|t-1} x_t' f_{t|t-1}^{-1} \tag{12}$$

is the so-called Kalman gain, which determines the weight assigned to new information about β_t contained in the prediction error.

The specification above is very general, underlying that some relevant aspects, like the betas' rule of motion, are far from obvious. That said, the results from our tests of structural changes on unconditional betas point to a particular dynamics for the time-varying parameters, and call for a specification that allows for multiple breaks. Indeed, Engle and Watson (1985) suggested to model as unit root processes the regression coefficients of relationships derived under the hypothesis that agents update their estimates only when new information becomes available, which is exactly our case. There is also a broad consensus in the literature (see for instance Santos and Veronesi, 2004) that betas are likely to be highly persistent quantities, essentially because they are functions of persistent shocks. Ang and Chen (2007) model betas as quasi-unit root $AR(1)$ processes. Adrian and Franzoni (2009) share the random walk assumption in their priors, whereas Fama and French (2006) do not. Lo and MacKinlay (1998) and Lo (2007) contain further empirical and theoretical support to the unit-root hypothesis. I also estimate my model positing an autocorrelation of 0.95 for conditional betas, obtaining no significantly different results. On the one hand, none of the TVK beta displays exploding behaviour (see below). On the other hand, some evidence shows that specifying the betas as explicit mean-reverting processes imposes costly restrictions on their dynamics. Indeed, this is what seems to happen in Jostova and Philipov (2005), where most estimated betas are not statistically different from unity and display very limited time variation. Hence, I assume that each regression coefficient in β_t follows a random walk, and in the following set $F = I_k$.

In equation (6), an inference on coefficients β_t given information up to time $t-1$ is a function of the inference on β_{t-1} given information up to time $t-1$, due to the law of motion of the state vector. Thus, uncertainty underlying $\beta_{t|t-1}$ is a function of the uncertainty underlying $\beta_{t-1|t-1}$ and Q, the covariance of the shocks to β_t. This is shown in equation (7).

More importantly, the prediction error in the time-varying parameter model consists of two parts: the component due to error in making an inference about β_t $\left(\text{i.e., } \beta_t - \beta_{t|t-1}\right)$, which is related to systematic risk, and the prediction error due to ε_t, the random shock to excess return R_t^e. Therefore, in equation (9), the conditional variance of the prediction error is a function of the uncertainty associated with $\beta_{t|t-1}$ and of S, the variance of ε_t. The updating equation in (10) suggests that $\beta_{t|t}$ is formed as a kind of weighted average of $\beta_{t|t-1}$ and new information contained in the prediction error $\eta_{t|t-1}$, the weight assigned to new information being the Kalman gain. Examining K_t more carefully, we notice that it is an inverse function of S, the variance of ε_t: the larger the idiosyncratic risk component, the smaller the weight assigned to new information about β_t contained in the prediction error. For a given market excess return in x_t, K_t is a positive function of the uncertainty surrounding $\beta_{t|t-1}$. On the other hand, if for simplicity we assume that β_t and x_t are 1×1, then the Kalman gain can be rewritten as

$$K_t = \frac{1}{x_t} \frac{P_{t|t-1} x_t^2}{P_{t|t-1} x_t^2 + S} \tag{13}$$

where $P_{t|t-1} x_t^2$ is the portion of the prediction error variance due to uncertainty in $\beta_{t|t-1}$ and S is the component due to the random shock ε_t. We can easily see that

$$\left| \frac{\partial K_t}{\partial \left(P_{t|t-1} x_t^2 \right)} \right| > 0 \tag{14}$$

suggesting that, when uncertainty associated with $\beta_{t|t-1}$ increases, relatively more weight is given to new information in the prediction error, $\eta_{t|t-1}$. Intuitively, the algorithm interprets an increase of uncertainty about $\beta_{t|t-1}$ as a deterioration of the information content of $\beta_{t|t-1}$ relative to that of $\eta_{t|t-1}$.

More complex approaches to the treatment of time variation and uncertainty in the market's assessment of risk are certainly viable. Ang and Chen (2007) employ a reduced-form version of the conditional CAPM in which only betas are allowed to change over time and conditional but constant alphas are extracted via numerical optimization. However, Ang and Chen's is a richly parametrized model (13 parameters vs. 5 in Jostova and Philipov and only 3 in the present study), and crucially relies on priors and assumptions about time variation in the mean and volatility of the conditional market return[16]. The KF methodology I adopt in this paper presents a number of advantages. First, it accounts for investors' uncertainty about asset risk in a straightforward way, as it entails a simple learning process on the model's coefficients: rational investors have to infer the factor loadings from observable portfolio returns and past prediction errors. The uncertainty they face depends upon the error variance of their past optimal forecast. Second, it is methodologically parsimonious, as its implementation requires narrow parametrization compared to, say, multi-equation settings, or alternative state-space models with regime-switching. Third, estimation is not based on conditioning information or strong assumptions about period-to-period variation in beta. For instance, this exercise does not employ restrictive assumptions as to the frequency of actual betas and their changes[17]. Fourth, it is consistent with a time-varying representation of CAPM in which uncertainty about current betas directly translates into changing conditional variance of returns. Finally, whereas most existing studies do not deal directly with the issue of parameter uncertainty and limited information on pricing errors, the approach in this paper endogenizes the realization of pricing errors and prevents future information from affecting today's forecasts.

5 Time variation and systematic risk

This Section reports estimates of TVK alphas, betas and other parameters from the one-factor model above. An additional useful feature of the KF specification chosen in this paper is that it yields endogenously consistent volatilities for alphas, betas, and idiosyncratic risk,

[16] Also, idiosyncratic risk is not allowed to change over time.

[17] Ozoguz (2009) considers a two-state regime-switching model of the aggregate market return, based on the idea that the economy fluctuates between two states. While interesting, this intuition however rests on a relatively rigid assumption on the number and nature of the states. This is why this exercise allows for smoother fluctuations of alphas and betas.

rather than forcing their computation through simulation or other approximations. Hence, I will comment on results beginning with the KF estimates of volatility parameters σ_α, σ_β and σ_ε for the 25 FF portfolios and their combinations HLS, $HL2$, $HL3$, $HL4$, HLB and SBH, $SB2$, $SB3$, $SB4$, SBL, reported in Table II (respective standard errors are in parentheses).

Several interesting findings emerge from Table II. First of all, alphas are very smooth. Their standard deviation is at most 0.06% per month, with extreme portfolio B1S5 and its combinations as exceptions. Such volatilities are on average much smaller than those implied by the OLS quarterly measures computed by Lewellen and Nagel (2006). This also says that the parameter variability allowed for by the KF algorithm does not introduce excess variation in the coefficients. Alphas of small-stock portfolios exhibit a slight tendency to higher volatility than those of large-stock portfolios, and the same applies to high BE/ME portfolios *vis-à-vis* low BE/ME ones. Both betas and idiosyncratic risk are definitely more volatile for small stocks than for large-stock portfolios; there is no discernible impact of BE/ME on this effect.

Matching the standard deviations with the time-varying estimates of alphas and betas reported below and the speed with which tje KF algorithm converges, we can conclude that TVK parameters are tightly estimated. It is also worth noting the smaller volatility of these coefficients compared with the constant-coefficient, OLS case that I presented above, and which does employ the same information set. This fits in well with findings by Ang and Chen (2007), who employ asymptotic theory to demonstrate that standard OLS inference provides misleading results, precisely because of time variation in the quantity of market risk. However, the magnitude and persistence of estimated parameters imply, as shown below, that TVK estimates lend much less support to CAPM's ability to explain the spreads in average returns than most results in Ang and Chen's study.

Figures I-IV plot, for some representative portfolios, the time series of interest, as generated by the procedure. The top panels of each figure show estimated conditional betas (*beta*) and, closer to the x-axis, alphas (*alpha*); the lower panels show the conditional forecast error variance of portfolio returns (*vareta*). Projected over the length of the estimation sample, and compared with the 60-month window often considered, most betas exhibit marked medium-term variation, typically over an interval of one year or two. Therefore, to discriminate between these fluctuations and more slow-moving drifts in betas, I also plot simple centered 30-month moving averages ($betaMA$)[18].

Let us start with portfolio B1S1, that is, the basket of highest BE/ME stocks based on the smallest firms. The dynamics of this portfolio's TVK parameters differs substantially from most other cases. Nevertheless, this portfolio's estimates reveal some properties of alphas and betas, like their quite different inertia, that will feature regularly in many other examples. Recall that over the 1926-2007 sample this portfolio commands the highest monthly average return amongst the 25 FF portfolios (1.74%), whereas its unconditional CAPM coefficients are $\alpha^u = 0.50\%$ and $\beta^u = 1.4$. Its TVK beta (Figure I) displays ample variation over the entire sample. We can neatly detect an upward trend between the late 1920s and mid-1940s (peaking to nearly 5 and with an average of 2.4), then a substantial fall for longer than a decade thereafter, some stabilization in the 1960s and then a resumption of the downward drift until early 2000s (trough about 0.7). The run-up to the 2007-09 financial crisis witnesses a sudden increase in the beta's value. Of equal interest, the TVK estimate of this portfolio's alpha steadily creeps up, starting from negative values in the 1930s and growing to or staying close around zero between the 1940s and the mid-1960s. Afterwards, it turns definitively positive, reaching and then overtaking its unconditional value towards the end of the sample[19].

Overall, the results for this portfolio, which is particularly skewed towards small-cap and high BE/ME stocks, reveal that CAPM's preferred measure of systematic risk trends downward over 1926-2007. However, substantial changes in both beta and alpha do take place over long sub-periods, with a gradual upward drift of alpha easily the clearest one. In contrast, we will see below that the estimated alphas of small-cap, low BE/ME portfolios are persistently negative over the whole sample. This arguably confirms that the conventional

[18] Alternatively, HP-filtering would be an equivalent smoothing transformation, albeit not more parsimonious.

[19] To avoid further cluttering the charts, I omit confidence bands -available upon request along with complete estimates. As Table II shows, all coefficients are precisely estimated.

CAPM fails to explain the so-called value effect, to a larger extent as the sample period gets longer.

In the lower chart of Figure I, I plot the conditional variance of the prediction error. The methodology yields *vareta* as a measure of uncertainty, and models it as endogenous with respect to both future idiosyncratic risk and the evolution of *beta*. For portfolio B1S1, we notice clusters of higher uncertainty in the 1930s and across the 1990s, whereas it was relatively subdued in 1940s-1970 and in the 2000s. Significantly, the considerable rise of this beta over 2001-2006 takes place alongside low conditional volatility. As a consequence of the model's structure, the largest absolute pricing errors occur in periods when uncertainty was greater. Results for portfolios B1S2, B1S3, B1S4 and B1S5, that is, the progressively larger-stock sorts with the highest BE/ME stocks, do not qualitatively differ from those for B1S1. This arguably means that market capitalization does not significantly alter the effects we uncover.

As a sample of findings for the low BE/ME sorts, let us now turn to B5S5, that is, the portfolio with lowest BE/ME stocks of the largest firms (average return 0.91%, $\alpha^u = -0.05\%$, $\beta^u = 0.98$). Volatilities in Table II say that this portfolio's beta varies within a much narrower band. Indeed, Figure II portraits a dynamics that differs markedly from what characterizes the small-value portfolio. Its TVK beta still tends to somehow rise and fall, but in this case never markedly departing from about unity, roughly its unconditional value. In contrast, its alpha follows a long downward trend from small, positive (and significant) values at the start of the sample up until small, negative (and significant) values in the mid-1980s. The growth in aggregate market values over 1983-2000 coincides with a reversal of this trend, whereas that in 2003-2007 witnesses a simultaneous fall of both alpha and beta. In sum, the beta of this low BE/ME, large-cap portfolio displays quite limited variation in the medium and long run, whereas alpha trends downwards over the whole sample. Here again, the largest pricing errors occur when the conditional variance of the forecast error exhibits larger values. However, the conditional variance of returns is comparatively small. The latter two findings further validate the choice made in Section 4 over how to measure uncertainty.

As for the smaller-cap, growth-stock portfolio, which we know to have on average a much smaller return, Figure III shows results for B5S2 (average return 0.88%; $\alpha^u = -0.26\%$, $\beta^u = 1.25$). Unlike the results for B5S5, same BE/ME category but larger firms, *alpha* here is anywhere negative and its magnitude is particularly sizeable in the 1930s and early 1940s. Its beta moves between 1 and 2, with large long-term swings, overall drifting upwards from the mid-1950s onwards. As with other small-cap portfolios, the years 1999 to 2006 witness a remarkable increase of its beta, despite its persistently negative or insignificant alpha. This is further evidence suggesting that sorting on BE/ME, much more than on market capitalization, is the main driver of alphas and of their evolution.

Finally, I chart estimates for portfolio B3S3 (average return 1.29%; $\alpha^u = 0.21\%$, $\beta^u = 1.15$), which lies exactly at the intersection of both BE/ME and size sorts, and therefore is not tilted towards any characteristic. Figure IV shows that its TVK beta reaches a maximum of about 1.5 in the 1930s, then gradually declines until the late 1950s, when it resumes an upward trajectory, culminating in the late 1960s. Since then, beta steadily falls in value, accelerating to a sharp drop around the dotcom bust, when it attains an all-time minimum of 0.5, suddenly reversed by a marked increase since then. This portfolio's alpha is always positive, although not sizeable, and persistent throughout the sample, except for the very early years. As for other cases, there are bouts of conditional uncertainty in the 1930s and around the 1990s.

5.1 Out-of-sample predictive performance

It is now vital to assess the relative predictive ability of the TVK and rolling OLS estimates, as the latter often feature in tests of conditional CAPM as well as in many, more practical uses. To this end, I evaluate the outcomes of a simple hedging strategy, implemented using either the rolling betas or the TVK ones. For each portfolio i, I compute the return over month $t+1$ of the position consisting of selling short β_t^j dollars of the market portfolio for each dollar invested in portfolio i, where β_t^j is time-t estimate of portfolio's beta obtained through either rolling OLS ($j = ROLS$) or the TVK algorithm ($j = TVK$)[20]. Formally, I calculate for each portfolio the return over month $t+1$ of the hedged long-short position, which is the hedging error h_{t+1}^{j}

$$h_{t+1}^{j} = R_{t+1} - \beta_t^j R_{t+1}^{eM},\ j = ROLS, TVK$$

For each month $t+1$, I estimate β_t^{ROLS} based on a rolling 60-month window ending in month t, whereas β_t^{TVK} is simply the month-t beta extracted via the KF technique. This simulation amounts to an out-of-sample test on the predictive ability of the TVK betas against the standard rolling-window OLS estimates often employed in the literature and the financial industry. Of course, the perfect hedge would yield a zero return. Table III reports for all portfolios the implied average hedging errors and their standard deviations.

Results are clear-cut. The absolute values of average hedging errors are smaller when β_t^{TVK} is used instead of β_t^{ROLS}. Only for 3 out of the 25 portfolios the opposite holds, and these are the "irregular" B1S5 portfolio, plus the two extreme small-cap, low BE/ME ones. Moreover, while almost all $avg\left(h_{t+1}^{ROLS}\right)$ are sizeable and negative, $avg\left(h_{t+1}^{TVK}\right)$ sligthly overshoots in some cases but it goes the other way round for the lowest BE/ME portfolios. This indicates that ROLS betas suffer from some bias that does not affect the TVK counterparts. Even more striking in terms of hedging efficiency, the ROLS model yields hedging errors that are always more volatile than with the rival TVK model[21]. This evidence supports unambiguously TVK betas against those from rolling-window regressions as forward-looking indicators of market-risk sensitivity.

6 The cyclical behavior of market-risk loadings

Are market betas tied to economic activity and market conditions? This is a long-standing question, whose relevance rises given the rich temporal and cross-sectional variation of the

[20] Jostova and Philipov (2005) perform a very similar exercise on industry portfolios.

[21] With the usual exception of portfolio B5S1. t-tests on the sample averages and volatilities (not shown here for brevity) show that those differences are almost always significant at conventional levels.

TVK estimates. To assess the interplay between time-varying betas and economic fluctuations, I perform two complementary exercises. First, I run simple regressions of each portfolio's TVK beta on a battery of state variables, at the monthly frequency. Second, I use quarterly averages of each monthly beta in unrestricted bivariate vector autoregressions (VARs) along with a quarterly business cycle indicator, to gauge the lead-lag relations among the variables at a more typical business-cycle frequency.

In the monthly regressions, the explanatory variables are: the value-weighted excess return on the market (MKT), the one-month Treasury bill rate (TBILL), the yield spread between ten-year and one-year Treasury bonds (TERM), the yield spread between Moody's seasoned Aaa and Baa corporate bonds (DEF), the log dividend yield on the value-weighted market index (DP), the consumption-to-wealth ratio of Lettau and Ludvigson (2001) (CAY)[22] and the log of the PMI Composite Index (PMI). To pin down the marginal power of the correlations with the dependent variable, I include all the variables jointly as regressors. Each of them is standardised, so that the resulting coefficient estimate can be interpreted as the change in TVK beta predicted by a one-standard-deviation change in the regressor. As before, computed standard errors are autocorrelation- and heteroskedasticity-consistent, following Andrews (1991)[23]. All these choices make the ensuing inference acceptably robust.

A few regularities emerge from the estimates, presented for each portfolio beta in Table IV. First, R^2s increase almost monotonically with the size of firms, reaching 54% for portfolio B2S5. State variables jointly explain a larger portion of time variation in the betas of larger firms' portfolios, particularly those characterized by high BE/ME, whereas those of small-stock portfolios are much more weakly related to them. Apparently, the market loadings of larger firms are more tightly associated with cyclical conditions. Second, the value-weighted excess return on the market and the Aaa-Baa spread seem to be almost orthogonal to betas, besides few correlations with the market loadings of some large-cap, low BE/ME portfolio. Third, CAY is the state variable most highly and systematically correlated with betas. CAY

[22] CAY is available only at the quarterly frequency. I computed monthly observations from the original data using linear interpolation.

[23] Due to data constraints, the estimation sample here starts in May 1953.

exhibits a strong and negative relationship with most betas, regardless of their BE/ME category. There are some exceptions among largest-cap portfolios, whereas for the smallest-cap portfolios this effect is stronger. Lettau and Ludvigson (2001) claim that their CAY indicator is broadly counter-cyclical. To the extent that the joint regressions I run do not dim this tendency, a negative coefficient would imply that betas evolve in a pro-cyclical way, particularly those of small-cap portfolios.

Fourth, when significant, a unit variation in TBILL has the largest impact on TVK betas among state variables. Regardless of market capitalization, the betas of high (low) BE/ME portfolios are negatively (positively) associated with TBILL. The latter is often found to have strong predictive power for economic activity (the same applying to TERM). This is additional, and more precise, evidence that the market-risk loadings of value (growth) portfolios tend to move with (against) the business cycle. It also confirms that high BE/ME portfolio betas are firmly procyclical, a finding partly supported by the positive associations of these portfolios' betas with PMI. Also, high (low) BE/ME betas are positively (negatively) correlated with DP, which implies that the betas of value stocks rise when the market valuation of cash flows declines. Lastly, the term spread does reveal some predictive content for TVK betas, but only for those of large-stock portfolios, whose betas have a positive correlation with the spread.

Taken together, these results say that some state variables, commonly used as leading indicators of the business cycle, also hold some useful information for developments in our betas. Betas of high BE/ME portfolios tend to move pro-cyclically, whereas those of low BE/ME portfolios lean towards an opposite tendency, albeit more weakly. Large-cap betas have stronger associations with state variables. These findings are particularly valuable, as the risk loadings were explicitly derived to account for the effects of uncertainty and time variation, but are based only on asset and market return data. It is appropriate to check whether these regularities are robust with respect to the choice of the data frequency, a different lead-lag structure, and alternative macroeconomic indicators. This is why next I examine the information content of TVK betas at a lower frequency and within a different statistical specification.

I compute simple three-month averages of each beta, and insert them in unrestricted bivariate VARs in which a TVK beta is jointly regressed on a quarterly business cycle indicator. Then, I run Wald tests on the significance of estimated coefficients to infer whether the indicator Granger-causes, i.e., is helpful in predicting future developments in betas, or rather the other way round. The logic of such exercise is straightforward. If markets process information efficiently, and if the TVK beta is an unbiased measure of the quantity of risk, its quarterly dynamics might be predicted by lagged indicators of business fluctuations. If instead, as broadly resulting from the monthly estimates, betas also contain coincident or forward-looking information about macroeconomic risk, its current value might help forecast future business conditions. Cross-sectionally, this exercise should also provide evidence as to whether sorting stocks according to BE/ME ratios and market capitalization leads to different patterns in the relationship between systematic risk and economic activity.

I employ three alternative indicators of aggregate fluctuations, all typically constructed at quarterly frequency: GDP-based output gap (YGAP, the deviation of actual from equilibrium real GDP), the log change in real personal consumption expenditure (PCE), and the log change in real private non-residential fixed investment (INV)[24]. I define the output gap by using the Congressional Budget Office's measure of potential GDP[25]. Given some data constraints, the estimation sample is from 1947Q2 (1949Q1 for YGAP) to 2007Q2. The VARs lag lengths are chosen according to conventional information criteria: this involves estimating a VAR(3), (4) or (5).

Table V lists the portfolios for which a TVK beta or a macroeconomic variable turns out to Granger-cause significantly (10% confidence level or less) one another. There are some clear regularities. Changes in aggregate consumption Granger-cause betas only weakly and only for few high BE/ME portfolios. On the other hand, betas do not predict future consumption growth. Therefore, consumption growth seems largely orthogonal to TVK betas.

[24] All data in this Section are from the Federal Reserve Bank of St. Louis' FRED database (http://research.stlouisfed.org/fred2)

[25] Alternative measures, obtained using Hodrick-Prescott and Baxter King filtering, yield very similar results.

The risk sensitivities of large-cap portfolios, both high and low BE/ME, tend to Granger-cause the output gap, whereas the latter does not help predict future betas. Finally, growth in aggregate investment predicts the market-risk loadings of high BE/ME and small firms portfolios. However, only betas of high BE/ME portfolios have some predictive content for future investment growth. Therefore, whilst in theory the cost of equity is a key determinant of a firm's decision to invest, it is investment growth that leads developments in TVK betas, rather than the other way round. The exceptions represented by some high BE/ME portfolios possibly reflect the fact that these sorts are skewed towards funding-constrained firms. These results lend some empirical support to the recent development of investment- and production-based asset pricing models (Liu et al., 2009).

7 Conclusions

The value of a firm reflects the sum of current investments plus the options to invest in all future projects (growth opportunities). In effect, there are no reasons to think that the exposures to fundamental sources of risk vary across assets, i.e. across claims to different cash flows, but not over time, i.e., when the information set and economic circumstances affecting the market valuation of the firm's cash flows likely change. Shifts in assets' unobserved risk loadings are likely to be the result of the arrival of new information as well as of investors adjusting their forecasts in a context of parameter and model uncertainty. Despite the growing interest of the finance literature about the role played by uncertainty and structural change, there are relatively few studies that investigate empirically their effects on investment choices.

I this paper, I employ a parsimonious learning model with no conditioning information to extract time-varying estimates of market-risk sensitivities, pricing errors and a measure of pricing uncertainty for the returns of portfolios sorted on book-to-market ratios and market capitalization. Estimated conditional parameters display significant fluctuations over time, along patterns that change across different size and BE/ME categories of stocks, although there is clear evidence that sorting on BE/ME, much more than on size, is the main driver of alphas and their evolution. The methodology I adopt allows for time variation in both the systematic and idiosyncratic components of stock returns, and yields time-varying betas that have markedly superior predictive ability for portfolio returns against rolling-window OLS estimates. I also examine the relationship of time-varying betas with state variables at business-cycle frequencies. Several other findings stand out. High (low) BE/ME betas move pro-cyclically (counter-cyclically). Large-cap portfolio betas have stronger associations with state variables. Finally, investment growth helps predict the betas of high BE/ME and small firms portfolios, whereas the betas of large-cap portfolios, both high and low BE/ME, have some predictive content for future output gap.

Looking ahead, this paper's approach, thanks to its simplicity and precision, lends itself easily to several developments. One application would extract TVK risk loadings in a multifactor framework. A further extension might involve a more formal test of conditional CAPM, in which the statistical significance of pricing errors is evaluated cross-sectionally for each sample observation. Finally, the method could also shed some light on the long-standing question of how conditional volatility enters the risk-return trade-off.

References

Adrian, T., and F. Franzoni (2009). Learning about Beta: time-varying factor loadings, expected returns and the conditional CAPM. Forthcoming, *Journal of Empirical Finance*.

Andersen, T.G., T. Bollerslev, F.X. Diebold, J. Wu (2005). A Framework for Exploring the Macroeconomic Determinants of Systematic Risk. *American Economic Review*, 95(2), 398-404.

Anderson, E.W., E. Ghysels, and J.L. Juergens (2009). The impact of risk and uncertainty on expected returns. Journal of Financial Economics, 94, 233-263.

Andrews, D.W.K (1991). Heteroskedasticity and Autocorrelation Consistent Covariance Matrix Estimation. *Econometrica*, 59, 817–858.

Andrews, D.W.K, and C.J. Monahan (1992). An Improved Heteroskedasticity and Autocorrelation Consistent Covariance Matrix Estimator. *Econometrica*, 60, 953–966.

Ang, A. and G. Beckaert (2007). Stock Return Predictability: Is it There? *Review of Financial Studies*, 20(3):651-707.

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

Bai, J., and Perron, P. (1998). Estimating and Testing Linear Models with Multiple Structural Changes. *Econometrica*, 66, 47-78.

Bai, J., and Perron, P., (2003). The Computation and Analysis of Multiple Structural Change Models. *Journal of Applied Econometrics*, 18, 1-22.

Bansal, R., R.F. Dittmar, and C.T. Lundblad (2005). Consumption, Dividends, and the Cross Section of Equity Returns. *Journal of Finance*, LX(4), 1639-1672.

Bekaert, G., E. Engstrom, and Y. Xing (2009). Risk, uncertainty, and asset prices. *Journal of Financial Economics*, 91, 59–82.

Campbell, J.Y., and T. Vuolteenaho (2004). Good beta, bad beta, *American Economic Review*, vol. 94(5), 1249–1275.

Cochrane, J.H. (2005). *Asset Pricing.* Princeton University Press.

Cooper, I., and R. Priestley (2009). Time-varying risk premiums and the output gap. *Review of Financial Studies*, 22:7.

Engle, R.F., and M.W. Watson (1985). Application of Kalman filtering in econometrics. World Congress of the Econometric Society, Cambridge, MA.

Fama, E. and K. French (1992a). The Cross-Section of Expected Stock Returns. *Journal of Finance*, 47(2), pp. 427-465.

Fama, E. and K. French (1992b). Industry cost of equity. *Journal of Financial Economics*, 43(2), 153-193.

Fama, E. and K. French (2006). The value premium and the CAPM. *Journal of Finance,* 61, 2163–2185.

Jagannathan, R., and Z. Wang (1996). The Conditional CAPM and the Cross-Section of Expected Returns. *Journal of Finance*, 51(1), 3-53.

Jostova, G. and A. Philipov (2005). Bayesian Analysis of Stochastic Betas. *Journal of Financial and Quantitative Analysis*, 40(4), 747-778.

Kim, C.-J. and C.R. Nelson (1999). *State-space models with regime switching.* MIT Press, Cambridge Mass.

Lettau, M., and S. Ludvigson (2001). Resurrecting the (C)CAPM: A Cross-Sectional Test When Risk Premia Are Time-Varying. *Journal of Political Economy*, 109(6).

Lewellen, J., and S. Nagel (2006). The conditional CAPM does not explain asset-pricing anomalies. *Journal of Financial Economics*, 82(2), 289-314.

Lewellen, J., S. Nagel and J. Shanken (2009). A skeptical appraisal of asset-pricing tests. *Journal of Financial Economics*, forthcoming.

Liu, L.X., T.M. Whited, and L. Zhang (2009). Investment-based expected stock returns. *Journal of Political Economy*, forthcoming.

Lo, A.W., and A.C. MacKinlay (1998). A non-random walk down Wall Street. Princeton University Press.

Lo, A.W. (2007). Efficient Markets Hypothesis. In L. Blume and S. Durlauf, eds., The New Palgrave: A Dictionary of Economics, Second Edition, 2007. New York: Palgrave Macmillan.

Mamaysky, H., M. Spiegel, and H. Zhang (2008). Estimating the Dynamics of Mutual Fund Alphas and Betas. *Review of Financial Studies*, 21, 233-264.

Ozoguz, A. (2009). Good Times or Bad Times? Investors' Uncertainty and Stock Returns. *Review of Financial Studies*, 22: 4377-4422.

Qu, Z., and P. Perron (2006). Estimating Restricted Structural Change Models, *Journal of Econometrics*, 134 373-399.

Pástor, L. and P. Veronesi (2009). Learning in financial markets. *Annual Review of Financial Economics*, Vol. 1.

Santos, T. and P. Veronesi (2004). Conditional betas. *NBER Working Paper Series* #10413.

Santos, T. and P. Veronesi (2009). Habit formation, the cross section of stock returns and the cash flow risk puzzle. *Journal of Financial Economics*, forthcoming.

Schwert, G.W. (1989a). Why does stock market volatility change over time? Journal of Finance, 44(5), 1115-1154.

Schwert, G.W. (1989b). Business Cycles, Financial Crises and Stock Volatility. *Carnegie-Rochester Conference Series on Public Policy*, 31, 83-125.

Zhang, L. (2005). The value premium. *Journal of Finance*, 60(1), 67-103.

Table I

Tests of multiple structural change on market betas: 1963-2007

The table contains the results of the tests proposed by Bai and Perron (1998, 2003) on the stability of estimated OLS coefficients for the intercept and slope of the linear regression $R_t^{ei} = a + bR_t^{eM} + e_t$,where $R_t^{ei} = R_t^i - R_t^f$ is the return on test portfolio i in excess of the one-month Treasury bill rate and $R_t^{eM} = R_t^{eM} - R_t^f$ is the excess return on the market. The $\sup \mathbf{F}_T(k)$ tests allow for the possibility of serial correlation in the disturbances. The HAC covariance matrix is constructed according to Andrews (1991) and Andrews and Monahan (1993) using a quadratic kernel with automatic bandwidth selection based on an $AR(1)$ approximation. The residuals are pre-whitened using a $VAR(1)$. Data are sampled at the monthly frequency, and span the period July 1963 to August 2007 (530 observations). *, **, *** indicate significance at the 90%, 95% and 99% level, respectively. The **WD** max statistics, when not significant, is displayed at the 5% significant level.

Test	**HLS**	**HL2**	**HL3**	**HL4**	**HLB**
$\sup \mathbf{F}_T(1)$	20.42***	6.98	3.24	15.01***	6.78
$\sup \mathbf{F}_T(2)$	10.59***	3.93	4.31	10.45***	6.25
$\sup \mathbf{F}_T(3)$	11.16***	5.95*	3.88	10.58***	14.65***
$\sup \mathbf{F}_T(4)$	8.82***	4.67*	4.13	10.74***	4.84*
$\sup \mathbf{F}_T(5)$	10.31***	4.15**	3.43**	8.78***	9.78***
UD max	20.42***	6.98	4.31	15.01***	14.65***
WD max	25.81***	8.41*	7.53	21.97***	24.48***
Test	**SBH**	**SB2**	**SB3**	**SB4**	**SBL**
$\sup \mathbf{F}_T(1)$	5.57	9.00**	16.46***	11.33**	12.89***
$\sup \mathbf{F}_T(2)$	3.78	3.80	9.62***	10.26***	7.45**
$\sup \mathbf{F}_T(3)$	3.60	6.98**	7.94***	9.04***	7.11**
$\sup \mathbf{F}_T(4)$	4.14	5.41**	6.72***	8.91***	6.02**
$\sup \mathbf{F}_T(5)$	2.36	3.31	5.00***	5.98***	3.85*
UD max	5.57	9.00**	16.46***	11.33**	12.89***
WD max	7.11	10.05**	16.46***	17.70***	12.89**

Table II
Volatilities of time-varying parameters for 25 size-BE/ME portfolios: 1926-2007

The table contains estimated time-varying parameter estimates from the model $R_t = x_t\beta_t + \varepsilon_t$, where $\beta_t = F\beta_{t-1} + \upsilon_t$, and σ_α, σ_β and σ_ε are conditional estimates of the standard deviations of one-factor regression coefficients and disturbance. R_t is the return on test portfolio i in excess of the one-month Treasury bill rate. Standard errors are in parentheses. Data are sampled at the monthly frequency and cover the period July 1926 to August 2007 (974 observations).

		Small		→		*Big*	
		B1S1	**B1S2**	**B1S3**	**B1S4**	**B1S5**	**SBH**
High	$\boldsymbol{\sigma}_\alpha$	0.02 (0.01)	0.02 (0.02)	0.03 (0.01)	0.01 (0.01)	4.86 (0.22)	2.39 (0.15)
	$\boldsymbol{\sigma}_\beta$	0.45 (0.03)	0.22 (0.03)	0.16 (0.02)	0.13 (0.02)	0.41 (0.04)	0.27 (0.03)
	$\boldsymbol{\sigma}_\varepsilon$	4.11 (0.12)	3.72 (0.12)	3.14 (0.08)	3.37 (0.09)	1.26 (0.44)	3.28 (0.16)
		B2S1	**B2S2**	**B2S3**	**B2S4**	**B2S5**	**SB2**
	$\boldsymbol{\sigma}_\alpha$	0.03 (0.02)	0.01 (0.01)	0.01 (0.01)	0.01 (0.01)	0.01 (0.01)	0.00 (0.01)
	$\boldsymbol{\sigma}_\beta$	0.41 (0.03)	0.19 (0.02)	0.11 (0.01)	0.13 (0.01)	0.06 (0.01)	0.31 (0.02)
	$\boldsymbol{\sigma}_\varepsilon$	3.37 (0.1)	2.86 (0.08)	2.45 (0.07)	2.17 (0.06)	2.43 (0.06)	2.76 (0.09)
		B3S1	**B3S2**	**B3S3**	**B3S4**	**B3S5**	**SB3**
↓	$\boldsymbol{\sigma}_\alpha$	0.00 (0.02)	0.03 (0.01)	0.00 (0.01)	0.01 (0.01)	0.01 (0.04)	0.00 (0.01)
	$\boldsymbol{\sigma}_\beta$	0.28 (0.03)	0.26 (0.02)	0.08 (0.01)	0.1 (0.01)	0.04 (0.01)	0.20 (0.02)
	$\boldsymbol{\sigma}_\varepsilon$	4.22 (0.12)	2.41 (0.07)	2.08 (0.05)	1.89 (0.05)	1.88 (0.05)	3.25 (0.1)
		B4S1	**B4S2**	**B4S3**	**B4S4**	**B4S5**	**SB4**
	$\boldsymbol{\sigma}_\alpha$	0.06 (0.03)	0.00 (0.01)	0.00 (0.01)	0.00 (0.00)	0.00 (0.02)	0.02 (0.02)
	$\boldsymbol{\sigma}_\beta$	0.69 (0.04)	0.26 (0.02)	0.06 (0.01)	0.06 (0.01)	0.04 (0.01)	0.48 (0.03)
	$\boldsymbol{\sigma}_\varepsilon$	4.54 (0.14)	2.82 (0.09)	2.23 (0.06)	1.71 (0.04)	1.38 (0.03)	4.00 (0.11)
		B5S1	**B5S2**	**B5S3**	**B5S4**	**B5S5**	**SBL**
Low	$\boldsymbol{\sigma}_\alpha$	0.03 (0.03)	0.00 (0.02)	0.00 (0.01)	0.00 (0.01)	0.00 (0.01)	0.02 (0.07)
	$\boldsymbol{\sigma}_\beta$	0.58 (0.06)	0.15 (0.03)	0.13 (0.02)	0.03 (0.01)	0.01 (0.00)	0.37 (0.04)
	$\boldsymbol{\sigma}_\varepsilon$	6.89 (0.21)	3.78 (0.12)	2.88 (0.09)	2.04 (0.05)	1.45 (0.03)	4.50 (0.14)
		HLS	**HL2**	**HL3**	**HL4**	**HLB**	
	$\boldsymbol{\sigma}_\alpha$	0.03 (0.02)	0.00 (0.01)	0.00 (0.01)	0.00 (0.01)	2.44 (0.14)	
	$\boldsymbol{\sigma}_\beta$	0.13 (0.03)	0.05 (0.01)	0.08 (0.01)	0.12(0.01)	0.17 (0.02)	
	$\boldsymbol{\sigma}_\varepsilon$	3.77 (0.12)	2.82 (0.07)	2.60 (0.07)	2.70 (0.07)	2.40 (0.14)	

Table III

This table contains the sample mean ($avg\left(h_{t+1}^{j}\right)$) and standard deviation ($\sigma\left(h_{t+1}^{TVK}\right)$) of the hedging error $h_{t+1}^{j}= R_{t+1}-\beta_t^j R_{t+1}^{eM}$ where the superscript $j = ROLS, TVK$ denotes the betas obtained either through a rolling 60-month window regression ending in month t, or via the Kalman algorithm. Data are monthly and the estimation sample is from 1928.6 to 2007.7.

	$avg\left(h_{t+1}^{ROLS}\right)$	$avg\left(h_{t+1}^{TVK}\right)$	$\sigma\left(h_{t+1}^{ROLS}\right)$	$\sigma\left(h_{t+1}^{TVK}\right)$
$B1S1$	-1.30	0.63	10.91	6.84
$B2S1$	-0.86	0.44	8.55	6.05
$B3S1$	-0.77	0.36	10.02	5.90
$B4S1$	-0.43	-0.21	11.66	7.48
$B5S1$	-0.20	-0.54	16.91	10.42
$B1S2$	-1.04	0.34	8.50	4.92
$B2S2$	-0.68	0.32	5.84	3.92
$B3S2$	-0.55	0.25	5.04	3.89
$B4S2$	-0.54	0.10	6.07	4.08
$B5S2$	-0.16	-0.27	7.16	4.58
$B1S3$	-0.85	0.27	8.83	4.13
$B2S3$	-0.49	0.25	4.00	2.93
$B3S3$	-0.42	0.23	3.82	2.45
$B4S3$	-0.35	0.19	3.32	2.49
$B5S3$	-0.20	-0.04	5.72	3.55
$B1S4$	-0.85	0.23	10.05	4.18
$B2S4$	-0.45	0.25	4.52	3.13
$B3S4$	-0.34	0.20	2.73	2.42
$B4S4$	-0.14	0.04	2.53	1.99
$B5S4$	-0.07	0.01	3.23	2.15
$B1S5$	-0.42	-1.24	6.84	13.92
$B2S5$	-0.16	0.00	4.80	2.71
$B3S5$	0.01	0.07	2.40	2.06
$B4S5$	0.09	0.00	1.65	1.53
$B5S5$	0.08	-0.08	1.57	1.53

Table IV

This table contains OLS estimates for the slope and R^2 of a regression of each portfolio's time-varying beta on one lag of all of the state variables together. *, **, *** indicate significance at the 90%, 95% and 99% level, respectively. HAC standard errors were computed, following Andrews (1991). Data are sampled at the monthly frequency, and the estimation sample is from May 1953 to August 2007 (652 observations).

	MKT	$TBILL$	$TERM$	DEF	DP	CAY	PMI	R^2
$B1S1$	0.00	−0.09*	−0.01	0.01	0.06	−0.12***	0.04	0.12
$B2S1$	0.01	−0.04	−0.01	0.00	0.05	−0.11***	0.03	0.09
$B3S1$	0.01	0.06*	0.02	−0.01	0.02	−0.12***	0.06**	0.14
$B4S1$	0.00	0.05	0.04	0.00	0.00	−0.13***	0.00	0.05
$B5S1$	−0.02	0.03	−0.06	0.06	−0.12**	−0.11**	−0.02	0.13
$B1S2$	0.00	−0.10**	0.00	0.01	0.09**	−0.08**	0.06*	0.18
$B2S2$	0.01	−0.1***	−0.05	0.04	0.04	−0.07**	0.03	0.12
$B3S2$	0.00	0.00	0.01	−0.03	0.05	−0.13***	0.02	0.14
$B4S2$	0.00	0.06*	0.04	0.01	0.00	−0.08**	0.04	0.08
$B5S2$	−0.01	0.18***	0.05*	0.00	−0.12***	−0.09***	0.02	0.28
$B1S3$	0.00	−0.10**	0.03	0.01	0.12***	−0.08***	0.08***	0.24
$B2S3$	0.01	−0.13***	−0.03	0.03	0.09***	−0.04*	0.06**	0.28
$B3S3$	0.00	−0.02	0.02	0.01	0.06**	−0.06***	0.03	0.19
$B4S3$	0.00	0.09***	0.05**	−0.02	0.00	−0.06***	0.03*	0.18
$B5S3$	−0.02*	0.20***	0.09***	−0.06*	−0.08***	−0.07**	−0.03	0.26
$B1S4$	0.01	−0.13***	−0.01	0.03	0.15***	−0.05*	0.09***	0.34
$B2S4$	0.01	−0.12***	−0.01	0.00	0.09***	−0.07***	0.04	0.25
$B3S4$	0.01	0.00	0.03	−0.02	0.06**	−0.02	0.04**	0.14
$B4S4$	0.00	0.08***	0.03**	0.02	0.02	−0.05***	0.04***	0.31
$B5S4$	−0.02***	0.10***	0.04***	0.01	−0.09***	−0.02*	0.00	0.44
$B1S5$	0.00	−0.12**	0.02	−0.06	0.11***	−0.07**	0.02	0.18
$B2S5$	0.02***	−0.10***	−0.02	−0.03**	0.14***	0.02	0.03*	0.54
$B3S5$	0.01**	0.00	0.01	0.01	0.03	0.00	0.04***	0.11
$B4S5$	0.00	0.03*	0.01	0.02	0.02	0.01	0.01	0.17
$B5S5$	0.00	−0.02***	0.00	−0.01	0.00	0.03***	0.00	0.34

Table V

This table contains the p-value of significant (at least 10% confidence level) χ^2statistic of Wald tests on estimated coefficients of unrestricted bivariate VARs of each portfolio's time-varying beta (quarterly average of original monthly estimates) and a macroeconomic indicator. HAC standard errors were computed (Andrews, 1991). Data are quarterly, and the estimation sample is from 1947Q2 (1949Q1 for YGAP) to 2007Q2.

$YGAP$		INV		PCE	
$B1S1$	$YGAP \to BETA: 0.86$ $BETA \to YGAP: 0.07$	$B1S1$	$INV \to BETA: 0.03$ $BETA \to INV: 0.05$	$B2S1$	$PCE \to BETA: 0.04$ $BETA \to PCE: 0.47$
$B2S3$	$YGAP \to BETA: 0.11$ $BETA \to YGAP: 0.02$	$B2S1$	$INV \to BETA: 0.01$ $BETA \to INV: 0.12$	$B1S2$	$PCE \to BETA: 0.07$ $BETA \to PCE: 0.54$
$B5S3$	$YGAP \to BETA: 0.40$ $BETA \to YGAP: 0.06$	$B3S1$	$INV \to BETA: 0.01$ $BETA \to INV: 0.11$	$B1S5$	$PCE \to BETA: 0.03$ $BETA \to PCE: 0.42$
$B1S4$	$YGAP \to BETA: 0.53$ $BETA \to YGAP: 0.00$	$B4S1$	$INV \to BETA: 0.04$ $BETA \to INV: 0.07$		
$B2S4$	$YGAP \to BETA: 0.79$ $BETA \to YGAP: 0.03$	$B1S2$	$INV \to BETA: 0.03$ $BETA \to INV: 0.08$		
$B5S4$	$YGAP \to BETA: 0.83$ $BETA \to YGAP: 0.00$	$B1S3$	$INV \to BETA: 0.18$ $BETA \to INV: 0.07$		
$B1S5$	$YGAP \to BETA: 0.17$ $BETA \to YGAP: 0.01$	$B2S3$	$INV \to BETA: 0.04$ $BETA \to INV: 0.12$		
$B2S5$	$YGAP \to BETA: 0.29$ $BETA \to YGAP: 0.01$	$B1S4$	$INV \to BETA: 0.07$ $BETA \to INV: 0.00$		
$B4S5$	$YGAP \to BETA: 0.44$ $BETA \to YGAP: 0.03$	$B5S4$	$INV \to BETA: 0.22$ $BETA \to INV: 0.09$		
$B5S5$	$YGAP \to BETA: 0.66$ $BETA \to YGAP: 0.09$	$B1S5$	$INV \to BETA: 0.03$ $BETA \to INV: 0.68$		
		$B2S5$	$INV \to BETA: 0.04$ $BETA \to INV: 0.15$		
		$B3S5$	$INV \to BETA: 0.52$ $BETA \to INV: 0.07$		

Figure I

Time-varying, conditional CAPM parameters: 1928-2007. Portfolio B1S1.

The top panel plots estimated time-varying parameters from the model $R_t^{ei} = x_t\beta_t^i + \varepsilon_t^i$, where $\beta_t^i = F^i\beta_{t-1}^i + \upsilon_t^i$ and x_t contains a constant and the market's excess return. b*etaMA* is the centered 30-month moving average of *beta*. The lower panel plots the estimated conditional variance $f_{t|t-1} = E\left[\eta_{t|t-1}^2\right]$ ($vareta$) of the prediction error $\eta_{t|t-1} = R_t^{ei} - R_{t|t-1}^{ei}$.

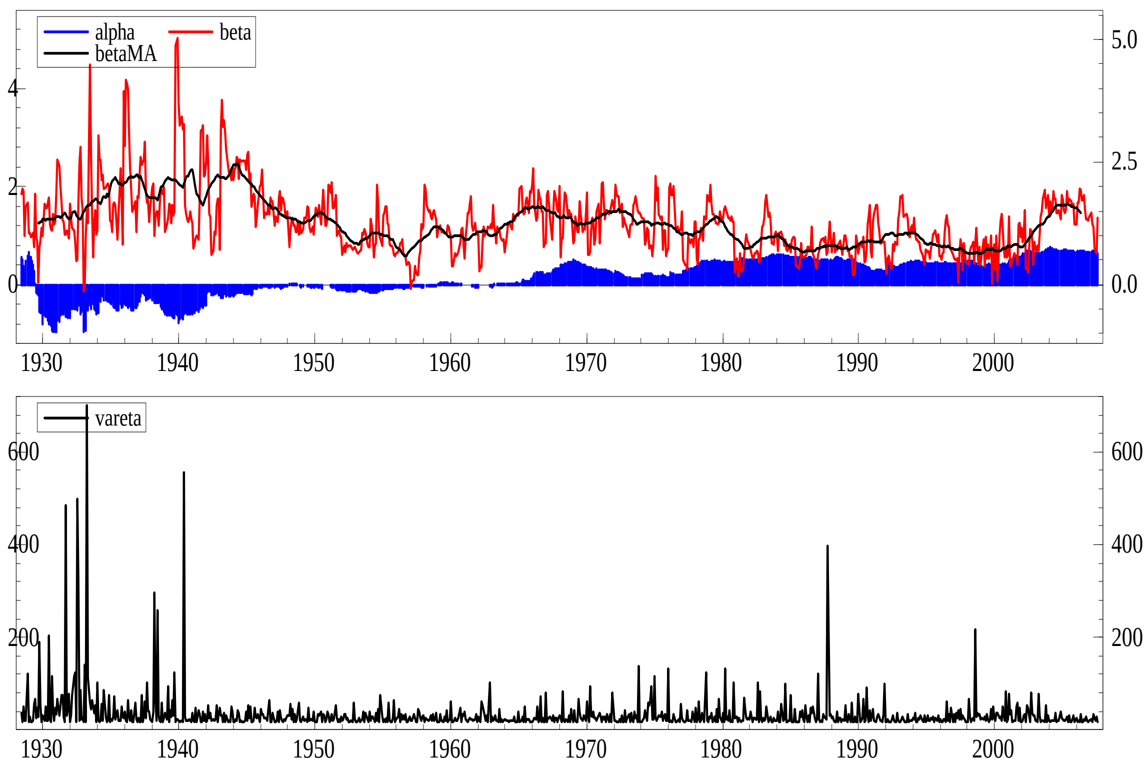

Figure II

Time-varying, conditional CAPM parameters: 1928-2007. Portfolio B5S5.

The top panel plots estimated time-varying parameters from the model $R_t^{ei} = x_t\beta_t^i + \varepsilon_t^i$, where $\beta_t^i = F^i\beta_{t-1}^i + \upsilon_t^i$ and x_t contains a constant and the market's excess return. b*etaMA* is the centered 30-month moving average of *beta*. The lower panel plots the estimated conditional variance $f_{t|t-1} = E\left[\eta_{t|t-1}^2\right]$ ($vareta$) of the prediction error $\eta_{t|t-1} = R_t^{ei} - R_{t|t-1}^{ei}$.

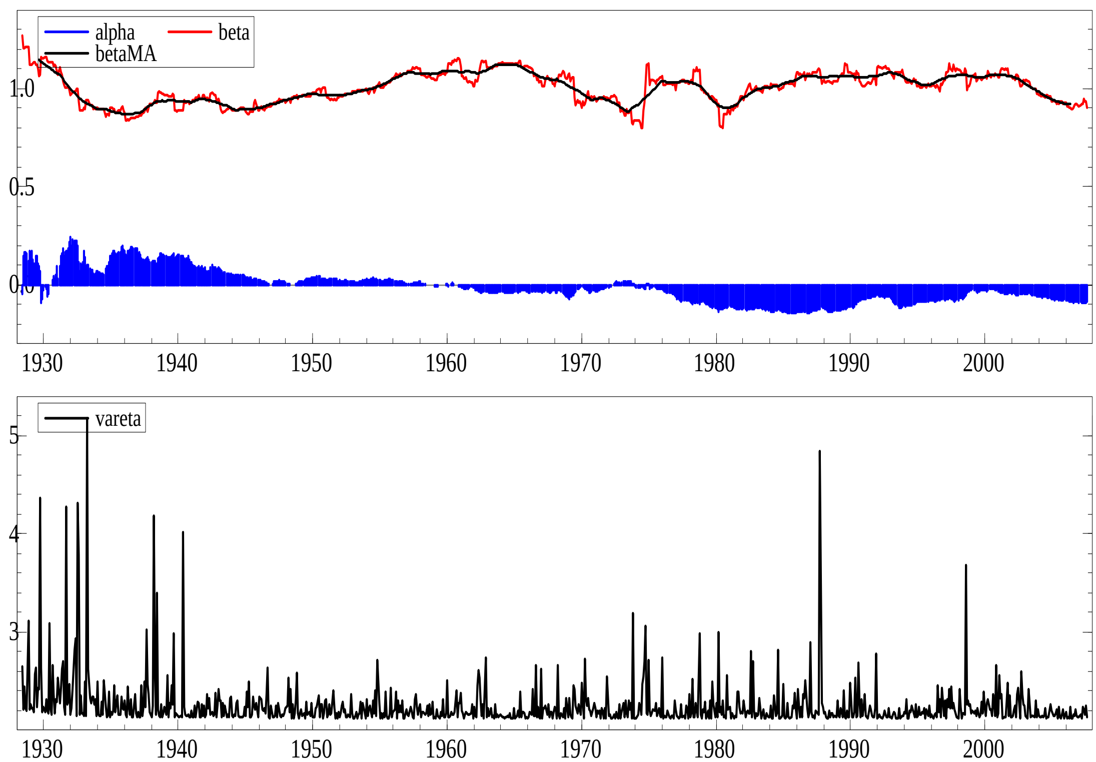

Figure III

Time-varying, conditional CAPM parameters: 1928-2007. Portfolio B5S2.

The top panel plots estimated time-varying parameters from the model $R_t^{ei} = x_t \beta_t^i + \varepsilon_t^i$, where $\beta_t^i = F^i \beta_{t-1}^i + \upsilon_t^i$ and x_t contains a constant and the market's excess return. b$etaMA$ is the centered 30-month moving average of $beta$. The lower panel plots the estimated conditional variance $f_{t|t-1} = E\left[\eta_{t|t-1}^2\right]$ ($vareta$) of the prediction error $\eta_{t|t-1} = R_t^{ei} - R_{t|t-1}^{ei}$.

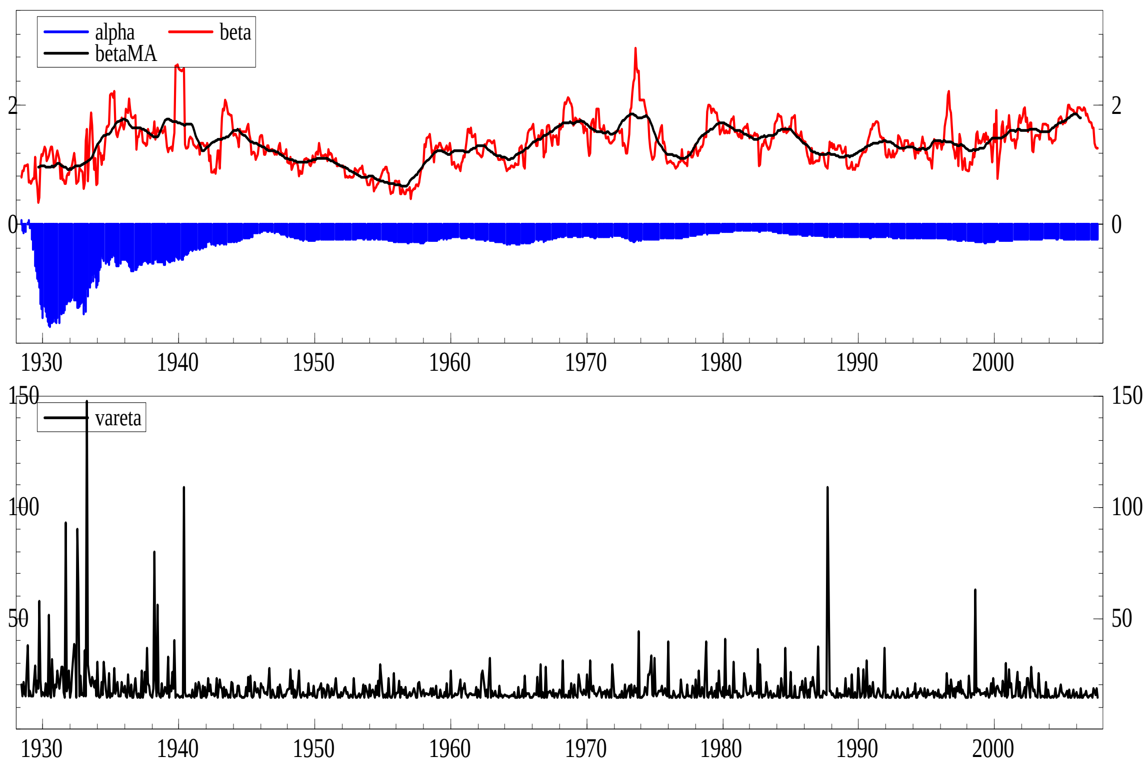

Figure IV

Time-varying, conditional CAPM parameters: 1928-2007. Portfolio B3S3.

The top panel plots estimated time-varying parameters from the model $R_t^{ei} = x_t \beta_t^i + \varepsilon_t^i$, where $\beta_t^i = F^i \beta_{t-1}^i + \upsilon_t^i$ and x_t contains a constant and the market's excess return. b*etaMA* is the centered 30-month moving average of *beta*. The lower panel plots the estimated conditional variance $f_{t|t-1} = E\left[\eta_{t|t-1}^2\right]$ ($vareta$) of the prediction error $\eta_{t|t-1} = R_t^{ei} - R_{t|t-1}^{ei}$.

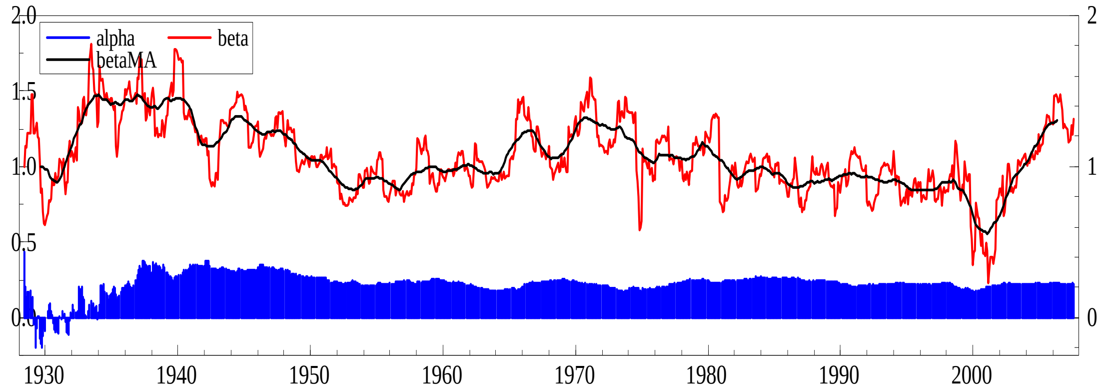

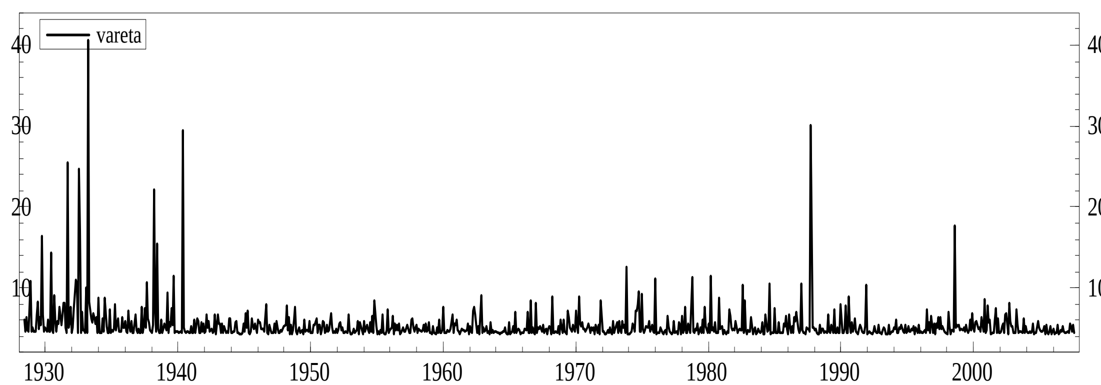

Appendix

Table A1 reports the average returns for the test assets and for the excess return on the market (R_t^{eM}), i.e., the value-weighted return on all NYSE, AMEX, and NASDAQ stocks minus the one-month Treasury bill rate (from Ibbotson Associates). Each portfolio is dubbed by its reference to the BE/ME and size category. In the notation, the firm size increases from S1 to S5 and the BE/ME ratio from B1 to B5. So, for instance, portfolio B2S5 contains stocks of the largest firms within the second highest BE/ME class, and so on. In addition to the 25 size-BE/ME portfolios, I report results for some of their combinations. $H - L$ is the difference between the average return for the two highest BE/ME quintiles within a size quintile and the average return of the two lowest BE/ME quintiles. Similarly, $S - B$ is the difference between the average return for the two smallest quintiles within a BE/ME quintile and the average return of the two largest quintiles within a BE/ME quintile. $\tau(\cdot)$ are the t statistics for the means of the time series, i.e., the average monthly return divided by its standard error.

There is a well-known data constraint. In the early part of the sample, some extreme portfolios often had very few firms. Limiting the sample to the period 1963 to 2007, as in Fama and French (2006) and Lewellen and Nagel (2006), allows considering portfolios with at least 10 stocks. The bottom panel of Table A1 contains statistics for this shorter sample. Aggregating over four portfolios for each dimension should make the inference more robust to the impact of extreme portfolios that are relatively under-diversified for the early part of the sample.

Next, I estimate for each portfolio return the standard one-factor model regression:

$$R_t^{ei} = a + bR_t^{eM} + e_t,$$

where $R_t^{ei} = R_t^i - R_t^f$ is the return on test portfolio i in excess of the one-month Treasury bill rate and again $R_t^{eM} = R_t^M - R_t^f$ is the excess return on the market. I estimate the regressions over both the whole sample (1926-2007) and the shorter period 1963-2007. Tables A2 and A3 below exhibit estimated coefficients a and b and their t-values[26].

Figure A1 helps gain some perspective on the extent to which conventional market betas fail to explain the so-called value and size premia. Following popular practice (Cochrane, 2005), the two panels plot excess returns against betas, for the subsample 1963-2007.

[26] t-values are constructed using HAC standard errors, following Andrews (1991).

Table A1

Average percent monthly returns for 25 size-BE/ME portfolios: 1926-2007 and 1963-2007

The table contains the returns on the FF portfolios, formed as the intersection of independent sorts of stocks on the basis of BE/ME ratios and market capitalization, and the difference between the value-weighted market return and the one-month Treasury bill rate. $H - L$ is the difference between the average return for the two highest BE/ME quintiles within a size quintile and the average return of the two lowest BE/ME quintiles. Similarly, $S - B$ is the difference between the average return for the two smallest quintiles within a BE/ME quintile and the average return of the two largest quintiles within a BE/ME quintile. $\tau(\cdot)$ are the t statistics for the means of the time series, i.e., the average monthly return divided by its standard error. Data are sampled at the monthly frequency, and cover the period July 1926 to August 2007 (top panel, 974 observations) and July 1963 to August 2007 (bottom panel, 530 observations).

	Small		$\rightarrow$		*Big*	$S-B$	$\tau(S-B)$
			Sample: 1926M7-2007M8				
High	B1S1	B1S2	B1S3	B1S4	B1S5	SBH	$\tau(SBH)$
	1.74	1.53	1.43	1.40	0.05	0.91	4.21
	B2S1	B2S2	B2S3	B2S4	B2S5	SB2	$\tau(SB2)$
	1.52	1.41	1.31	1.28	1.05	0.30	2.52
$\downarrow$	B3S1	B3S2	B3S3	B3S4	B3S5	SB3	$\tau(SB3)$
	1.35	1.35	1.29	1.19	1.00	0.26	1.92
	B4S1	B4S2	B4S3	B4S4	B4S5	SB4	$\tau(SB4)$
	1.13	1.26	1.20	1.06	0.92	0.20	1.17
Low	B5S1	B5S2	B5S3	B5S4	B5S5	SBL	$\tau(SBL)$
	0.78	0.88	0.99	0.99	0.91	-0.12	-0.62
$H-L$	HLS	HL2	HL3	HL4	HLB	R_t^{eM}= 0.66	
	0.67	0.40	0.28	0.31	-0.37		
$\tau(H-L)$	$\tau\,(HLS)$	$\tau\,(HL2)$	$\tau\,(HL3)$	$\tau\,(HL4)$	$\tau\,(HLB)$	$\tau\left(R_t^{eM}\right) = 3.84$	
	4.79	3.81	2.64	2.47	-1.71		
			Sample: 1963M7-2007M8				
High	B1S1	B1S2	B1S3	B1S4	B1S5	SBH	$\tau(SBH)$
	1.44	1.50	1.49	1.36	1.09	0.35	2.41
	B2S1	B2S2	B2S3	B2S4	B2S5	SB2	$\tau(SB2)$
	1.52	1.43	1.33	1.33	1.06	0.28	2.12
$\downarrow$	B3S1	B3S2	B3S3	B3S4	B3S5	SB3	$\tau(SB3)$
	1.32	1.40	1.22	1.21	0.96	0.27	1.86
	B4S1	B4S2	B4S3	B4S4	B4S5	SB4	$\tau(SB4)$
	1.28	1.15	1.22	1.00	0.97	0.23	1.39
Low	B5S1	B5S2	B5S3	B5S4	B5S5	SBL	$\tau(SBL)$
	0.70	0.87	0.91	1.00	0.88	-0.15	-0.86
$H-L$	HLS	HL2	HL3	HL4	HLB	R_t^{eM}= 0.49	
	0.59	0.46	0.35	0.34	0.15		
$\tau(H-L)$	$\tau\,(HLS)$	$\tau\,(HL2)$	$\tau\,(HL3)$	$\tau\,(HL4)$	$\tau\,(HLB)$	$\tau\left(R_t^{eM}\right) = 2.61$	
	4.32	3.58	2.62	2.75	1.25		

Table A2

Unconditional one-factor model coefficients for 25 size-BE/ME portfolios: 1926-2007

The table reports OLS estimates for the intercept, slope and R^2 of the linear regression $R_t^{ei} = a + bR_t^{eM} + e_t$, where $R_t^{ei} = R_t^i - R_t^f$ is the return on test portfolio i in excess of the one-month Treasury bill rate and $R_t^{eM} = R_t^M - R_t^f$ is the excess return on the market. t-values are in parentheses and are constructed using HAC standard errors, following Andrews (1991). Data are sampled at the monthly frequency, and cover the period July 1926 to August 2007 (974 observations).

	B1S1	**B1S2**	**B1S3**	**B1S4**	**B1S5**	**SBH**
a	0.50(3.04)	0.32(2.21)	0.20(1.38)	0.13(0.91)	-1.08(-1.26)	0.58(1.31)
b	1.40(15.8)	1.37(17.4)	1.40(14.7)	1.45(14.1)	1.25(12.5)	0.04(0.66)
R^2	0.61	0.71	0.76	0.75	0.25	0.00
	B2S1	**B2S2**	**B2S3**	**B2S4**	**B2S5**	**SB2**
a	0.34(2.27)	0.29(2.41)	0.26(2.61)	0.19(1.86)	-0.01(-0.13)	-0.08(-0.67)
b	1.31(16.9)	1.23(18.5)	1.13(21.9)	1.81(17)	1.14(13.5)	0.11(2.15)
R^2	0.66	0.76	0.80	0.82	0.79	0.03
	B3S1	**B3S2**	**B3S3**	**B3S4**	**B3S5**	**SB3**
a	0.12(0.74)	0.26(2.43)	0.21(2.41)	0.16(1.98)	0.04(0.46)	-0.22(-1.78)
b	1.40(17)	1.19(17.3)	1.15(21.5)	1.09(28.5)	0.98(20.9)	0.26(4.13)
R^2	0.66	0.76	0.85	0.87	0.85	0.11
	B4S1	**B4S2**	**B4S3**	**B4S4**	**B4S5**	**SB4**
a	-0.16(-0.95)	0.11(0.92)	0.14(1.74)	0.02(0.27)	0.00(0.02)	-0.34(-2.86)
b	1.48(13.8)	1.28(20.2)	1.13(35.4)	1.10(31.8)	0.93(50.1)	0.37(4.84)
R^2	0.55	0.76	0.86	0.90	0.91	0.12
	B5S1	**B5S2**	**B5S3**	**B5S4**	**B5S5**	**SBL**
a	-0.62(-3)	-0.26(-1.90)	-0.17(-1.59)	-0.03(-0.33)	-0.05(-0.86)	-0.71(-4.83)
b	1.66(15.6)	1.25(19.9)	1.29(31.9)	1.08(28)	0.98(60.8)	0.43(6.92)
R^2	0.51	0.71	0.81	0.86	0.92	0.14
	HLS	**HL2**	**HL3**	**HL4**	**HLB**	
a	0.51(4.04)	0.07(0.69)	-0.07(-0.57)	-0.14(-1.15)	-0.83(-1.87)	
b	-0.21(-4.27)	0.04(0.56)	0.06(0.90)	0.23(2.43)	0.24(3.12)	
R^2	0.06	0.00	0.01	0.10	0.04	

Table A3

Unconditional one-factor model coefficients for 25 size-BE/ME portfolios: 1963-2007

The table reports OLS estimates for the intercept, slope and R^2 of the linear regression $R_t^{ei} = a + bR_t^{eM} + e_t$,where $R_t^{ei} = R_t^i - R_t^f$ is the return on test portfolio i in excess of the one-month Treasury bill rate and $R_t^{eM} = R_t^M - R_t^f$ is the excess return on the market. t-values are in parentheses, and are constructed using HAC standard errors, following Andrews (1991). Data are sampled at the monthly frequency, and cover the period July 1963 to August 2007 (530 observations).

	B1S1	**B1S2**	**B1S3**	**B1S4**	**B1S5**	**SBH**
a	0.66(3.52)	0.51(3.12)	0.53(3.17)	0.40(2.82)	0.20(1.50)	-0.18(-1.29)
b	1.03(19.2)	1.06(19.8)	1.00(18.4)	0.99(19.6)	0.83(17.4)	0.13(3.31)
R^2	0.58	0.66	0.67	0.68	0.58	0.03
	B2S1	**B2S2**	**B2S3**	**B2S4**	**B2S5**	**SB2**
a	0.56(3.23)	0.48(3.42)	0.41(3.15)	0.40(3.35)	0.20(1.77)	-0.25(-1.84)
b	1.00(19.9)	0.98(22.2)	0.91(22.1)	0.92(23.3)	0.80(21.9)	0.13(3.45)
R^2	0.61	0.70	0.72	0.75	0.80	0.03
	B3S1	**B3S2**	**B3S3**	**B3S4**	**B3S5**	**SB3**
a	0.31(1.80)	0.41(2.95)	0.28(2.17)	0.26(2.32)	0.07(0.73)	-0.27(-1.90)
b	1.08(21.8)	1.04(22.9)	0.98(23.3)	0.98(26)	0.85(28.2)	0.14(3.76)
R^2	0.63	0.72	0.77	0.80	0.77	0.03
	B4S1	**B4S2**	**B4S3**	**B4S4**	**B4S5**	**SB4**
a	0.20(1.02)	0.10(0.72)	0.20(1.65)	0.00(0.00)	0.03(0.36)	-0.33(-2.06)
b	1.23(24)	1.17(26.8)	1.11(33.5)	1.08(29.8)	0.95(39.8)	0.19(4.71)
R^2	0.60	0.74	0.81	0.85	0.87	0.05
	B5S1	**B5S2**	**B5S3**	**B5S4**	**B5S5**	**SBL**
a	-0.48(-1.99)	-0.31(-1.84)	-0.24(-1.77)	-0.09(-0.71)	-0.09(-1.08)	-0.78(-4.26)
b	1.45(25.9)	1.44(32.1)	1.36(35.9)	1.26(42.8)	1.01(45.7)	0.32(6.55)
R^2	0.61	0.73	0.78	0.84	0.88	0.11
	HLS	**HL2**	**HL3**	**HL4**	**HLB**	
a	0.28(1.93)	0.13(0.95)	0.02(0.12)	-0.03(-0.22)	-0.24(-1.93)	
b	-0.32(-7.81)	-0.28(-8.58)	-0.28(-6.37)	-0.20(-4.56)	-0.16(-4.11)	
R^2	0.20	0.18	0.16	0.1	0.07	

Figure A1
Average percent excess returns against unconditional market betas: 1963-2007

Top panel: lines connect portfolios with different BE/ME categories within size categories. Bottom panel: lines connect portfolios with different size categories within BE/ME categories.