Netspar   Network for Studies on Pensions, Aging and Retirement

Rob Bauer

Mathijs Cosemans

Rik Frehen

Peter Schotman

A Bayesian Panel Data Approach to Explaining Market Beta Dynamics

A Bayesian Panel Data Approach to Explaining Market Beta Dynamics

Rob Bauer[a,b], Mathijs Cosemans[a*], Rik Frehen[a,b], and Peter Schotman[a,b,c]

March 12, 2008

Abstract

We characterize the process that drives the market betas of individual stocks by setting up a hierarchical Bayesian panel data model that allows a flexible specification for beta. We show that combining the parametric relationship between betas and conditioning variables specified by economic theory with the robustness of an autoregressive specification delivers superior estimates of firm-specific betas. Our model also improves the accuracy of beta forecasts, which we use to construct optimal portfolios subject to target beta constraints. We further provide empirical support for the prediction of conditional asset pricing theory that individual stocks exhibit different risk dynamics. Finally, we document strong cross-sectional heterogeneity in firm-specific betas within the 25 size-B/M portfolios that are commonly used to test asset pricing models.

*Corresponding author: Mathijs Cosemans, Maastricht University, P.O. Box 616, 6200 MD Maastricht, Netherlands, e-mail: M.Cosemans@finance.unimaas.nl.

Introduction

Precise estimates of firm-specific betas are crucial in many applications of modern finance theory, including asset pricing, corporate cost-of-capital calculations, risk management, and performance evaluation. For instance, portfolio managers often have to ensure that their market risk exposure stays within predetermined limits and managers need accurate estimates of a company's beta to make optimal capital budgeting decisions. However, existing approaches to measuring market betas either assume that all stocks within a given portfolio have the same beta (e.g., Fama and MacBeth (1973)) or yield imprecise estimates of stock-specific betas because they estimate separate time-series regressions for each firm (e.g, Brennan, Chordia, and Subrahmanyam (1998)). In addition, although a large body of empirical evidence indicates that betas vary over time, existing studies do not provide clear guidance how market beta dynamics should be modeled.

Our main objective in this paper is to obtain more accurate estimates of the market risk of individual stocks by combining the benefits of the two main approaches to modeling variation in betas.[1] The first one, proposed by Shanken (1990), models conditional betas by allowing them to depend linearly on a set of conditioning variables. Rational asset pricing theory posits that the predictive power of firm characteristics and macroeconomic variables for stock returns is due to their relation with risk. Gomes, Kogan, and Zhang (2003) set up a theoretical framework in which size and B/M are correlated with the true conditional market beta. Santos and Veronesi (2004) show within a general equilibrium model that market betas vary substantially with the state of the economy. Empirical evidence that systematic risk is related to firm characteristics and business cycle variables is provided by, among others, Jagannathan and Wang (1996), Lewellen (1999), Ferson and Harvey (1999), Lettau and Ludvigson (2001), and Avramov and Chordia (2006). While economically appealing, the main drawback of this approach is that the investor's set of conditioning information is unobservable. Ghysels (1998) shows that misspecifying beta risk may result in serious pricing errors that might even be larger than those produced by an unconditional asset pricing model. In addition, this method produces excessive variation in betas due to sudden spikes in macroeconomic variables.

---

[1]An alternative approach has recently been proposed by Christoffersen, Jacobs, and Vainberg (2007), who calculate forward-looking betas using the information embedded in option data. A drawback of this method is that it requires a cross-section of liquid individual stock options, which is not available for many small firms.

The second, non-parametric, approach to model risk dynamics is based on purely data-driven filters. Approaches in this category include estimating short-window regressions (Lewellen and Nagel (2006)), rolling regressions (Fama and French (1997), Hoberg and Welch (2007)), and modeling beta as a latent autoregressive process estimated using the Kalman filter (Jostova and Philipov (2005), Ang and Chen (2007)). While these methods preclude the need to specify conditioning variables it is not clear which factors explain the cross-sectional and time variation in market beta. In addition, the time series of conditional betas obtained by following a purely data-driven approach will lag the true variation in beta, because using a window of past returns to estimate the beta at time t will give an estimate of the average beta during this time window.

Our panel data approach uses both high-frequency return data and cross-sectional firm-level data to capture the time-series dynamics and cross-sectional heterogeneity in beta. Combining time-series and cross-sectional information produces superior estimates of firm-specific betas because previous studies document a strong cross-sectional relationship between beta and firm characteristics and because Bollerslev and Zhang (2003) show that the use of high-frequency data yields more precise and timelier estimates of beta than those produced by a rolling window approach. We implement the MIDAS approach developed by Ghysels, Santa-Clara, and Valkanov (2005) to estimate realized betas, which differs from traditional rolling window estimators of betas by mixing data sampled at different frequencies and by choosing the optimal weights given to past data.

We estimate the model using a sample of 5,017 NYSE-AMEX stocks over the period July 1964 through December 2006. One of the virtues of using a large panel of individual stocks is that it allows us to study the cross-sectional characteristics of conditional alphas and betas. Using individual stocks instead of the commonly used characteristics sorted portfolios offers several other advantages. First, portfolios conceal individual security characteristics that can be important determinants of stock returns. Second, it is more interesting to study the betas of individual firms because these exhibit much more time variation than portfolio betas. A third important drawback of the portfolio approach is that only a few characteristics can be addressed simultaneously, because otherwise many portfolios would consist of only a few stocks.

However, despite these benefits the use of individual stocks is still limited compared to the use of characteristics sorted portfolios. An important reason for this is that although the cross-

section of firms is often quite large, the time-series dimension of most data sets is rather short. Due to this small time-series it is difficult to estimate the model parameters with a reasonable degree of precision. Brennan, Chordia, and Subrahmanyam (1998) use a sample of individual stocks to test static asset pricing models and Avramov and Chordia (2006) extend this approach to study the performance of conditional asset pricing models. Because these authors estimate a separate time series regression for every firm, their beta estimates are noisy.

We contribute to these studies by estimating a Bayesian panel data model that exploits the large cross-section of firms to obtain more precise parameter estimates. In particular, we model heterogeneity in parameters by specifying hierarchical priors. Thus, we impose a structure on coefficients by assuming that they are drawn from a common distribution. This enables us to capture the cross-sectional variation in market beta without the need to estimate a large number of parameters. Intuitively, as explained by Hsiao and Pesaran (2007), the Bayes estimator can be interpreted as a weighted average of the least squares estimator for a given cross-section unit and the cross-sectional average coefficient. Specifically, the Bayes estimator of the firm-specific parameters shrinks the least squares estimator towards the cross-sectional mean. When the number of time series observations for a firm increases, the weight gradually shifts from the prior to the data. Another benefit of our Bayesian methodology is that it allows for exact inference, avoiding the need to rely on asymptotic distributions, which can lead to severe small sample problems as documented by Ang and Chen (2007).

We find that modeling stock-specific betas as a function of both conditioning variables and realized betas dominates traditional specifications in which betas depend on conditioning variables or realized betas alone. Combining these specifications produces superior beta estimates because they capture different aspects of market beta dynamics. Moreover, in contrast to traditional rolling window estimators of beta that give equal weight to past returns, we show that it is preferable to give more weight to recent observations. We also show that our model improves the accuracy of stock-specific beta forecasts, which we use to construct optimal portfolios subject to target beta constraints. In addition, we provide empirical support for the prediction of conditional asset pricing theory that individual stocks exhibit significantly different risk dynamics. Specifically, we find that the cross-sectional variation in beta increases sharply during recessionary periods. Finally, we document strong cross-sectional heterogeneity in firm-specific

3

betas within the 25 size-B/M portfolios that are commonly used to test asset pricing models. This violation of the homogeneity assumption underlying the portfolio approach has important implications for empirical tests of asset pricing models.

The paper proceeds as follows. In section 1 we introduce our specification for modeling firm-specific conditional betas in a panel data framework. Section 2 discusses our choice of prior parameter distributions and explains the Bayesian approach to inference. Section 3 describes the data set. We report our empirical results in section 4. Section 5 concludes.

1   The Model

In this section we describe our conditional beta specification. We set up a hierarchical Bayesian panel data model that allows a flexible specification for market beta. We start from a model for excess returns on individual stocks,

$$r_{it} = \alpha_{it} + \beta_{it} r_{Mt} + \epsilon_{it}, \tag{1}$$

where r_{it} is the excess return on stock i in month t, α_{it} is the risk-adjusted return and represents model mispricing, β_{it} is the conditional market beta, r_{Mt} is the excess market return, and ϵ_{it} is a zero-mean, normally distributed idiosyncratic return shock.

Our specification for the conditional market beta consists of two components: one part is the past realized beta, b_{it-1}, and the other part is the past fundamentals-based beta, β_{it-1}^*,

$$\beta_{it} = \phi_i b_{it-1} + (1 - \phi_i)\beta_{it-1}^* + \eta_{it}, \tag{2}$$

where ϕ_i and $(1 - \phi_i)$ measure the proportion of the beta of firm i that is explained by the past realized beta and fundamentals-based beta, respectively, and where η_{it} is a zero-mean, normally distributed idiosyncratic shock to beta.

b_{it-1} is the past realized beta that we estimate using daily data according to the Mixed Data Sampling (MIDAS) approach introduced by Ghysels, Santa-Clara, and Valkanov (2005). Bollerslev and Zhang (2003) show that the use of high-frequency data yields more precise estimates of realized betas than those produced by a rolling regressions approach. We choose to estimate realized betas using daily returns because these provide a reasonable balance between

efficiency and robustness to microstructure noise. For instance, using intraday returns would lead to biased estimates of realized betas for less liquid stocks, due to non-synchronous trading effects (see, e.g., Campbell, Lo, and MacKinlay (1997)). The MIDAS approach differs from traditional rolling window estimators of betas by mixing data sampled at different frequencies and by selecting the optimal window for estimating betas using a flexible weighting function. Ghysels, Santa-Clara, and Valkanov (2005) use the MIDAS approach to estimate the market's conditional variance and find that it is superior to traditional GARCH and rolling window methods because of the extra power that comes from the use of mixed-frequency data and from the flexible form for the weights on past data.

In particular, our MIDAS estimator of realized betas is given by:

$$b_{it} = \frac{\sum_{\tau=1}^{\tau^{max}} w_{t-\tau} r_{it-\tau}^{(d)} r_{Mt-\tau}^{(d)}}{\sum_{\tau=1}^{\tau^{max}} w_{t-\tau} r_{Mt-\tau}^{(d)} r_{Mt-\tau}^{(d)}}, \tag{3}$$

where t refers to a particular month, τ to a particular trading day, and $w_{t-\tau}$ to the weight given to the product of the return on stock i and the market return, $r_{it-\tau}^{(d)} r_{Mt-\tau}^{(d)}$, and to the squared market return, $r_{Mt-\tau}^{(d)} r_{Mt-\tau}^{(d)}$, on day $t - \tau$. We set the maximum window length τ^{max} equal to 250 days, which is approximately one year of trading days. We parameterize the weights as a beta function:

$$w_{t-\tau} = \frac{f\left(\frac{\tau}{\tau^{max}}, \kappa_1; \kappa_2\right)}{\sum_{\tau=1}^{\tau_{max}} f\left(\frac{\tau}{\tau^{max}}, \kappa_1; \kappa_2\right)}, \tag{4}$$

where $f(\frac{\tau}{\tau^{max}}, \kappa_1; \kappa_2)$ is the density of a beta distribution.

As pointed out by Ghysels, Santa-Clara, and Valkanov (2005), the specification based on the beta function has several advantages. First, it ensures that the weights are positive and sum to one. Second, it is parsimonious because only two parameters need to be estimated. Third, it is flexible as it can take various shapes for different values of the two parameters. We impose a downward sloping pattern on the weights by setting κ_1 equal to 1, which further reduces the number of parameters that need to be estimated. $\kappa_1 = \kappa_2 = 1$ implies equal weights, which corresponds to a rolling window estimator of beta on daily data. $\kappa_1 = 1$ and $\kappa_2 > 1$ correspond to the case of slowly decaying weights. In general, the higher κ_2, the faster the rate of decay.

β_{it-1}^* is the fundamentals-based beta, parameterized as a function of conditioning variables,

$$\beta_{it-1}^* = \delta_0 + \delta_1'[Z_{it-1} \otimes BC_{t-1}], \tag{5}$$

where Z_{it-1} is a vector that contains L firm characteristics and BC_{t-1} is a vector that contains a constant and M business cycle variables. This specification allows the relation between beta and firm characteristics to vary over the business cycle. Modeling beta dynamics as a linear function of a set of predetermined instruments goes back to Shanken (1990) and is consistent with the economic motivation for conditional asset pricing models, in which the stochastic discount factor is a function of macroeconomic state variables and factor premia.

We model conditional alphas as a linear function of the same set of instruments:

$$\alpha_{it-1} = \alpha_{0i} + \alpha_1'[Z_{it-1} \otimes BC_{t-1}], \tag{6}$$

where α_{0i} is an individual effect that measures time-invariant mispricing.

We include both firm-specific and macroeconomic variables as instruments for conditional alphas and betas because of their documented predictive power for returns (Fama and French (1989) and Lewellen (1999)). Empirical evidence that systematic risk is related to firm characteristics and business cycle variables is provided by, among others, Jagannathan and Wang (1996), Lettau and Ludvigson (2001), and Avramov and Chordia (2006). The theoretical motivation for choosing firm characteristics as instruments is given by Gomes, Kogan, and Zhang (2003), who show that the ability of size and book-to-market to explain the cross-section of returns is due to their correlation with the true conditional market beta. Apart from size and B/M, we also select firm-specific momentum and turnover as conditioning variables to examine whether momentum and turnover effects are related to beta dynamics. Theoretical support for including macroeconomic variables is provided by Santos and Veronesi (2004), who show within a general equilibrium model that market betas vary substantially with the business cycle. Our choice of business cycle variables is motivated by previous work (e.g., Ferson and Harvey (1999)) and includes the default spread, dividend yield, one-month T-bill rate, and term spread.

Substituting equations (2), (5), and (6) into equation (1) leads to the following specification:

$$r_{it} = \alpha_{0i} + \alpha_1'[Z_{it-1} \otimes BC_{t-1}] + \phi_i b_{it-1} r_{Mt} + (1 - \phi_i)(\delta_0 + \delta_1'[Z_{it-1} \otimes BC_{t-1}])r_{Mt} + \eta_{it} r_{Mt} + \epsilon_{it}.$$

(7)

Our main objective in this paper is to determine whether the time series dynamics and cross-sectional variation in betas is better explained by firm characteristics and macroeconomic state variables, by realized betas, or by a linear combination of both. Therefore, we are primarily interested in the parameter ϕ_i and compare three different specifications based on equation (7): (1) Conditional beta (2) Fundamentals-based beta ($\phi_i = 0$) (3) Realized beta ($\phi_i = 1$).

2 Methodology

2.1 Prior Distributions

We specify conditionally conjugate, hierarchical priors that impose a common structure on the model parameters while still allowing parameters to vary across firms. Thus, our setup combines the benefits of a portfolio approach to estimating betas (e.g., Fama and MacBeth (1973)) and an approach in which separate regressions are estimated for each firm (e.g., Avramov and Chordia (2006)). We use relatively uninformative (diffuse) priors to minimize their influence on the posterior densities. Following Jostova and Philipov (2005), we specify noninformative prior distributions for the variance parameters $\sigma_{a_0}^2$, σ_ϕ^2, $\sigma_{\delta_0}^2$, $\sigma_{\epsilon_i}^2$, and σ_η^2 by setting the scale and shape parameters A and B of their inverse gamma (IG) prior distributions equal to 0.001. We set the degrees of freedom parameters ψ of the Wishart priors for $\Omega_{\alpha_1}^{-1}$ and $\Omega_{\delta_1}^{-1}$ equal to the dimensions of these matrices, because these values give the lowest possible weight to the prior information (Gelman, Carlin, Stern, and Rubin (2004)). We set the scale matrices S of the Wishart priors equal to the identity matrix. Furthermore, we give equal prior weight to the fundamentals-based beta and the realized beta by setting the prior mean of ϕ_i equal to 0.5.[2]

---

[2]We also considered specifications with prior mean of ϕ_i set equal to 0 or 1. Our results are robust to the choice of this prior distribution.

Specifically, our choice of prior distributions is as follows:

$$\alpha_{0i} \sim N(0, \sigma_{\alpha_0}^2) \qquad \text{with} \quad \sigma_{a_0}^2 \sim IG\left(0.001, 0.001\right),$$

$$\alpha_1 \sim N(0, \Omega_{\alpha_1}) \qquad \text{with} \quad \Omega_{\alpha_1}^{-1} \sim Wish\left([(L+LM)I]^{-1}, (L+LM)\right),$$

$$\phi_i \sim N(0.5, \sigma_\phi^2) \qquad \text{with} \quad \sigma_\phi^2 \sim IG\left(0.001, 0.001\right),$$

$$\delta_0 \sim N(0, \sigma_{\delta_0}^2) \qquad \text{with} \quad \sigma_{\delta_0}^2 \sim IG\left(0.001, 0.001\right),$$

$$\delta_1 \sim N(0, \Omega_{\delta_1}) \qquad \text{with} \quad \Omega_{\delta_1}^{-1} \sim Wish\left([(L+LM)I]^{-1}, (L+LM)\right),$$

$$\sigma_{\epsilon_i}^2 \sim IG(0.001, 0.001),$$

$$\sigma_\eta^2 \sim IG(0.001, 0.001).$$

We parameterize the MIDAS weights as a beta function and set κ_1 equal to 1. To rule out cases where more recent data receives less weight than observations in the more distant past, i.e., when $\kappa_2 < 1$, we constrain κ_2 to the interval [1,26]. When $\kappa_2 = 1$ all 250 days receive equal weight in the estimation and when $\kappa_2 = 26$ the cumulative weight given to the 40 most recent days is 99%. We implement this restriction by a change of variable, $\kappa_2 = 1 + 25\kappa_2^*$. For κ_2^* we choose a uniform prior, $\kappa_2^* \sim U[0,1]$. Because the conditional posterior density of κ_2 has a nonstandard form, we cannot directly sample from it. Therefore, we use the Metropolis-Hastings algorithm, in which candidate parameter values are drawn from a proposal density and accepted with a certain probability that is highest in areas of the parameter space where the posterior density is highest (see Chib and Greenberg (1995)). Details are provided in the appendix, which also presents the derivation of the joint posterior density and the conditional posterior distributions.

2.2 Bayesian Inference

We employ Markov Chain Monte Carlo (MCMC) methods to sample from the joint posterior distribution of the parameter vector θ. The main idea is to construct a Markov chain such that the chain converges to a unique stationary distribution that is the posterior density, $p(\theta|y)$. We use a particular MCMC algorithm, the Gibbs sampler, which involves the sequential drawing from the full conditional posterior densities to obtain draws from the joint posterior density. In particular, first the parameter vector θ is partitioned into B blocks $(\theta^{(1)}, \theta^{(2)}, ..., \theta^{(B)})$. At each

iteration of the Gibbs sampler each block is sampled from its posterior distribution conditional on all other blocks and the data. After a given number of iterations the chain converges and the draws are equivalent to draws from the joint posterior. We check convergence by inspecting the standardized cumsum statistics, suggested by Bauwens, Lubrano, and Richard (1999), applying the partial means test based on numerical standard errors, explained by Geweke (2005), and calculating the Gelman-Rubin statistic that compares the variation in output between and within chains, described by Gelman, Carlin, Stern, and Rubin (2004). These diagnostics indicate that the parameter chains have converged after 1,000 iterations. In our empirical analysis we therefore run 5,000 iterations and discard the first 1,000 iterations as burn-in period. The remaining draws are used to summarize the posterior density and to conduct Bayesian inference.

3 Data

The firm data comes from CRSP and Compustat and consists of the monthly return, size, book-to-market value, and turnover for a sample of NYSE- and AMEX-listed stocks. To calculate high-frequency-based realized betas we further retrieve daily returns from CRSP. The sample covers the period from July 1964 to December 2006. Following Avramov and Chordia (2006), we include a stock in the analysis for a given month t if it satisfies the following criteria. First, its return in the current month t and in the previous 36 months has to be available. Second, data should be available in month t-1 for size as measured by market capitalization, for the book-to-market ratio, and for turnover. We calculate the book-to-market ratio using accounting data from Compustat as of December of the previous year. Finally, in line with Fama and French (1993), we exclude firms with negative book-to-market equity. Imposing these restrictions leaves a total 5,017 stocks over the full sample period and an average of 1,815 stocks per month.

Table 1 presents summary statistics for the data set. Panel A reports the mean, median, standard deviation and 5th, 25th, 75th, and 95th percentile values of excess stock returns and firm characteristics across all data points. The average monthly excess stock return is 0.69% while the median is -0.16%. The mean (median) firm size is $1.59 (0.16) billion. Because the book-to-market ratio contains some extreme values, we decide to trim all book-to-market outliers to the 0.5th percentile and 99.5th percentile values of the distribution. After trimming, the average (median) book-to-market ratio equals 0.96 (0.75). Average turnover is 5.19% and

median turnover is 3.06% per month. The cumulative return over the twelve months prior to the current month, which we use as a proxy for momentum, has a mean of 14.65% and a median of 8.60%. Because the distributions of the firm characteristics display considerable skewness, we use their logarithmic transformations in the analysis. Furthermore, we normalize the characteristics by expressing them as deviations from their cross-sectional means to remove any time trend in the average value of the characteristics.

We further retrieve data for the four macroeconomic variables that we use as instruments for the conditional alphas and betas, i.e., the default spread, dividend yield, one-month Treasury bill rate, and term spread. We define the default spread as the yield differential between bonds rated BAA by Moody's and bonds with a Moody's rating of AAA. The dividend yield is calculated as the sum of the dividends paid on the value-weighted CRSP index over the previous 12 months divided by the current level of the index. The term spread is defined as the yield difference between ten-year and one-year Treasury bonds. Panel B shows descriptive statistics for the macroeconomic variables. The average default spread is 1.02%, the mean dividend yield equals 3.01%, the average one-month T-bill rate is 5.69%, and the average term spread is 0.85%.

[Table 1 about here.]

4 Market Beta Dynamics

In section 4.1 we examine whether betas are driven by lagged conditioning variables or by past realized betas. Section 4.2 compares the forecasting power of our conditional beta model to that of other specifications. Section 4.3 reports cross-sectional characteristics of alphas and betas and section 4.4 focuses on the time series dynamics of conditional alphas and betas. Section 4.5 discusses the plausibility of the homogeneity assumption underlying the portfolio approach to testing asset pricing models.

4.1 Fundamental Beta versus Realized Beta

A key objective in this paper is to characterize the process that governs the market betas of individual stocks. We investigate whether the time-series and cross-sectional variation in conditional betas is best explained by lagged firm characteristics and macroeconomic variables, by past realized betas, or by a linear combination of both. We address this question by estimating

10

the model in equation (7) and examining the posterior distributions of ϕ_i, which measures the proportion of beta explained by past realized beta. Figure 1 shows the cross-sectional distribution of the posterior means of the ϕ_i. In this figure we give more weight to the ϕ_i that are estimated with higher precision by weighting each posterior mean by its posterior precision. The cross-sectional average of the posterior mean is 0.54, which implies that both firm characteristics and macroeconomic variables as well as past realized betas are needed to accurately model conditional market betas. However, the cross-sectional standard deviation of the posterior means is large (0.36), suggesting that for some firms past realized betas are more important for explaining betas while for others lagged fundamentals-based betas have a stronger impact on betas. In sum, Figure 1 indicates that the specification in which beta dynamics are modeled as a function of both conditioning variables and realized betas dominates specifications in which betas depend on conditioning variables or realized betas alone.

[Figure 1 about here.]

Since we find that economically-motivated conditioning variables are important determinants of market betas, we now consider the posterior distributions of the parameters of the fundamentals-based beta. Table 2 presents the mean, median, standard deviation, and 5th, 25th, 75th, and 95th percentile values of the posterior distribution of the δ_0 and δ_1 parameters. The constant term (δ_0), which can be interpreted as the average fundamentals-based beta, has a posterior mean of 1.15 with a standard deviation of 0.02. The most important determinant of fundamental betas is book-to-market. Firm size has a negative impact on beta whereas book-to-market is positively related to beta. The finding that firms with a smaller market capitalization and higher B/M ratio have higher market betas is consistent with the theoretical work of Gomes, Kogan, and Zhang (2003). In particular, they demonstrate that size captures the component of a firm's systematic risk related to its growth options whereas the book-to-market ratio is a measure of the risk of the firm's assets in place. Interestingly, Table 2 further shows that momentum is positively related to market risk, suggesting that past winners are more risky than losers. In contrast, turnover does not drive variation in betas. The posterior distributions of the coefficients on the interaction terms between firm characteristics and macroeconomic variables indicate that the default spread and term spread are the most important business cycle variables for capturing variation in market betas.

11

[Table 2 about here.]

An important feature of our model is that it uses the MIDAS approach of Ghysels, Santa-Clara, and Valkanov (2005) to estimate realized betas based on daily return data. This approach incorporates a flexible weighting function that makes it possible to choose the optimal weights given to past data in estimating realized betas. The optimal window strikes a balance between giving equal weight to observations to obtain more precise beta estimates and giving more weight to recent data to obtain betas that are timelier and therefore more relevant. As shown in equation (4), we use a beta weighting function whose shape is determined by two parameters, κ_1, which we set equal to 1, and κ_2, which we estimate using the Metropolis-Hastings algorithm. $\kappa_1 = \kappa_2 = 1$ implies equal weights, which corresponds to the traditional rolling window estimator of beta. $\kappa_1 = 1$ and $\kappa_2 > 1$ correspond to the case of slowly decaying weights.

We find that in our conditional beta specification the posterior mean (standard deviation) of κ_2 is equal to 1.16 (0.04). Figure 2 compares the optimal weighting scheme implied by the posterior mean of κ_2 to the equal weighting scheme used by rolling window estimators. The plot shows that in the optimal scheme the most recent 150 days receive more weight than in the equal weighting scheme whereas the 100 most distant daily returns receive less weight. Thus, it is preferable to attach the highest weight to recent returns because these are most informative for estimating realized betas.

[Figure 2 about here.]

We now turn to the conditional betas produced by our model. We first calculate at each iteration of the Gibbs sampler the beta for firm i at time t. Subsequently, we compute the time-series average of these conditional betas. We then calculate for each firm the posterior mean of its time-series average beta. We give more weight to betas that are estimated with higher precision by weighting each posterior mean by its posterior precision. Figure 3 shows the cross-sectional distribution of these precision-weighted average firm betas. The distribution is centered around one and has a standard deviation of 0.34.[3] A 90% confidence interval for beta ranges from 0.46 to 1.60, which implies that firms differ substantially in their sensitivity to broad market movements.

---

[3]The cross-sectional average of betas deviates slightly from one because we weight firms by posterior precision and because our sample consists of NYSE-AMEX firms whereas the market index also includes NASDAQ stocks.

[Figure 3 about here.]

We also observe large cross-sectional differences in conditional alphas. Our calculation of average firm-specific alphas follows the procedure outlined above for average betas. Figure 4 shows the cross-sectional distribution of precision-weighted average firm alphas. The cross-sectional average of alpha is equal to 0.08, which implies that on average the model does a reasonable job in explaining returns. However, the cross-sectional variation in alphas is considerable, which indicates that for some firms the model leaves large pricing errors.

[Figure 4 about here.]

4.2 Beta Forecasts and Optimal Portfolios

In this section we compare the forecasting power of our conditional beta specification to that of the realized beta, fundamental beta, and static beta specifications. We assess the accuracy of the firm-specific beta forecasts produced by each of these models by constructing well diversified portfolios subject to target beta constraints.[4] This application illustrates the practical value of accurate beta estimates for portfolio management and risk management purposes.

At the end of each month we construct optimal portfolios that minimize the sum of squared portfolio weights subject to the constraint that the one-month ahead forecast of the portfolio beta is equal to the target beta. Thus, the portfolio construction at the end of month t-1 uses the forecasts available at time t-1 of the individual betas at time t. We allow short-selling of stocks but impose the constraint that portfolio weights have to sum to one, i.e., that the portfolio is fully invested. We then calculate the time-series of portfolio returns using the optimal weights. We perform this exercise for two different target portfolio betas, $\beta_p = 1$ and $\beta_p = 0$. The first target corresponds to the situation of a portfolio manager who has to ensure that his portfolio tracks the market index in terms of systematic risk. The second objective is relevant for hedge funds that follow a market neutral strategy and need accurate forecasts of firm-specific betas to neutralize their portfolio's market beta.

We measure the forecasting power of the four alternative approaches to modeling betas in two ways. First, we regress the time-series of portfolio returns on the market index and examine

---

[4]This application is inspired by Ghysels and Jacquier (2006), who minimize the ex-ante residual portfolio variance subject to market beta constraints. However, they focus on quarterly beta forecasts, optimize using industry portfolios rather than individual stocks, and only compare the forecasting power of data-driven specifications of conditional betas.

whether the realized portfolio beta equals the target beta. Second, we include interaction terms between the market index and the one-month ahead forecasts of the portfolio beta produced by the other three specifications in the regression and test whether these forecasts have explanatory power. For example, if we construct the optimal portfolio using the beta forecasts of the realized beta model, we add the beta forecasts for this portfolio generated by the conditional beta, fundamental beta, and static beta specifications. If a model is superior to alternative models, deviations from the target beta should not be predictable by these other beta specifications. Thus, our tests are based on the following time-series regression,

$$r_{pt} = \alpha_p + \beta_p r_{Mt} + \gamma_p' X_{pt-1} r_{Mt} + \epsilon_{pt}, \tag{8}$$

where X_{pt-1} is the vector that includes the one-month ahead forecasts of the portfolio beta produced by the specifications other than the one used to construct the portfolio. Our first test corresponds to the hypothesis that $\beta_p = \beta_{target}$. The second test corresponds to the hypothesis that the parameter vector $\gamma_p' = 0$.

Table 3 reports the results of this exercise. Panel A shows the regression output for the portfolios with a target beta of one. The estimates show that for all four specifications the realized portfolio beta does not significantly differ from the target beta of one. The realized beta of the portfolio that is based on the beta forecasts of the conditional beta specification is exactly equal to one. The realized portfolio betas corresponding to the fundamental beta and realized beta specifications equal 0.99. The realized beta of the portfolio that is constructed using static betas is 1.06, indicating that the static CAPM leads to the largest deviation (bias) from the target beta. To test our second hypothesis, we add the forecasts of the portfolio beta produced by the other specifications to the regression, orthogonalized with respect to each other and with respect to the beta forecast produced by the model used to construct the portfolio. The results in panel A show that the only portfolio for which all other beta forecasts are insignificant is the portfolio constructed using the beta forecasts of the conditional beta model. For the portfolio formed using the fundamental beta specification, the forecasts of the conditional and realized beta models are significant at the 1% level. The portfolio based on the realized beta model has significant exposures to the beta forecasts of the conditional beta model and the fundamental beta model. Finally, for the portfolio that is constructed using static betas

14

the forecasts of all other three specifications are significant.

These results indicate that the firm-specific beta forecasts produced by the conditional beta model are superior to those of the realized beta, fundamental beta, and static beta models. It appears that the realized beta model misses the information contained in the firm-specific and macroeconomic conditioning variables while the fundamental beta model misses the information incorporated in past returns. This supports our finding in the previous section that both specifications capture different aspects of market beta dynamics and that combining the two produces more accurate estimates of firm-specific betas.

The results for the portfolios with a target beta of zero are reported in panel B. For all four models the realized betas display a significant upward bias. The smallest bias is produced by the conditional beta model ($\beta_p = 0.02$) and the largest bias is generated by the realized beta model ($\beta_p = 0.07$). This upward bias in realized betas for portfolios with a target beta equal to zero is consistent with the results of Ghysels and Jacquier (2006). They attribute the systematic overshooting of the target beta to estimation errors in betas. Specifically, because the average beta across all stocks is one, stocks with the lowest betas receive the largest weight in the optimization in order to set the portfolio beta equal to zero. Ghysels and Jacquier (2006) argue that the smallest beta estimates contain negative estimation errors. Because the optimization overweights these stocks, the realized beta of the portfolio will be positive.

Panel B further shows that for the portfolio constructed using the forecasts of the conditional beta model the deviation from the target beta is predictable by the fundamental beta specification. In turn, the portfolio based on the fundamental beta model loads significantly on the forecasts of both the conditional beta model and the static beta model. The portfolios formed using the forecasts of the realized beta and static beta models have significant exposures to the conditional beta and fundamental beta specifications. Because the portfolio constructed using the forecasts of the conditional beta specification exhibits the smallest deviation from the target beta and has a significant exposure to only one of the other beta forecasts, we conclude that the conditional beta model is also the preferred specification for target betas equal to zero.

[Table 3 about here.]

15

4.3 Cross-Sectional Characteristics of Alphas and Betas

One of the virtues of using a large panel of individual stocks is that it allows us to carefully study the cross-sectional characteristics of conditional alphas and betas. In section 4.1 we documented considerable cross-sectional variation in the time-series means of alphas and betas. In this section we examine how the cross-sectional variation in firm-specific alphas and betas and the correlation between alpha and beta evolve over time. We relate these movements to the state of the economy.

Figure 5 shows the evolution through time of the cross-sectional standard deviation of conditional betas. The shaded areas in the plot indicate NBER recession periods. The average cross-sectional standard deviation is 0.43. However, the figure shows that the cross-sectional variation in market betas exhibits wide swings. An interesting result is that at the end of every recession period the variation in betas increases sharply. The most striking increase is in the beginning of the eighties where the cross-sectional variation almost doubled. This pattern is consistent with the economic motivation for conditional asset pricing models. In particular, conditional asset pricing theory asserts that firms differ in the sensitivity of their conditional betas to business cycle variations (see, e.g., Petkova and Zhang (2005) and Zhang (2005)).

[Figure 5 about here.]

Figure 6 compares the cross-sectional characteristics of the betas produced by our conditional beta specification to those of the betas obtained from the models in which betas depend on lagged conditioning variables or past realized betas alone. As a benchmark we also present results for the unconditional CAPM, in which betas are constant. The figure shows that the average cross-sectional variation in beta is largest in the model that only includes past realized betas, followed by the specification including both fundamental and realized betas, and the model that only includes fundamentals-based betas. The static CAPM produces the smallest cross-sectional variation in betas.[5]

However, the specification that only includes conditioning variables exhibits the largest swings in the cross-sectional variation in market risk, especially during recessionary periods. The high responsiveness of the fundamentals-based beta to macroeconomic conditions implies

---

[5]The cross-sectional variation in static betas varies over time only because firms enter and leave the sample.

16

that the conditioning variables capture important heterogeneity in risk dynamics across stocks. In contrast, the cross-sectional variation in realized betas is less responsive to business cycle variations. Moreover, it appears to react with a lag of a few months, for instance after the 1987 stock market crash. This lag arises because using a window of past returns to estimate realized betas gives an estimate of the average beta during this period. As expected, the time-series behavior of the cross-sectional variation in betas produced by the conditional beta specification in equation (2) is a combination of the dynamics of these two specifications. Thus, it combines the benefits of both specifications, responding fast to changes in economic conditions without producing too extreme variations in beta. Finally, due to its static nature the cross-sectional variation in market betas in the unconditional CAPM does not respond at all to business cycle variations. As a result, the static CAPM fails to capture changes in the cross-sectional distribution of stock returns.

[Figure 6 about here.]

In Figure 7 we plot the cross-sectional variation in conditional alphas. The average cross-sectional standard deviation is 2.65% but again we document a strong, positive relation between the state of the economy and the cross-sectional standard deviation of alphas. During the oil crisis in the seventies the cross-sectional standard deviation increases from 2% to more than 6.5% and during the recession in the 1980s the standard deviation reaches 7.5%. We attribute the sharp increase in the cross-sectional variation in alphas during crisis periods to the ability of some firms to successfully deal with the circumstances while others go bankrupt. Put differently, a recession separates the wheat from the chaff. Interestingly, the variation in alphas gradually declines towards the end of our sample period, suggesting that our model is better able to explain the cross-sectional variation in stock returns in recent times.

[Figure 7 about here.]

Figure 8 links together the cross-sectional dynamics in conditional alphas and betas by showing the evolution through time of the cross-sectional correlation between alphas and betas. The average cross-sectional correlation is slightly negative (-0.13). However, during recession periods this correlation becomes much more negative, reaching -0.8 in the beginning of the eighties. Thus, in times of economic distress stocks with the highest market beta tend to earn the lowest

17

returns, even after accounting for market risk. Note that towards the end of the sample period the correlation gradually increases to a level around zero.

[Figure 8 about here.]

4.4 Time Series Dynamics of Conditional Alphas and Betas

The previous section documents that the cross-sectional variation in alphas and betas and the correlation between alpha and beta vary strongly over time, particularly during recessions. We now examine the time series dynamics of alphas and betas by calculating for every firm the time series standard deviation and first-order autocorrelation of alpha and beta.

Panel A in Table 4 reports the cross-sectional distribution of the time-series variation in alpha. The specification that only includes past realized betas leads to the lowest time-series standard deviation in pricing errors, whereas the model that only includes fundamentals-based betas generates the highest standard deviation. In the conditional beta model of equation (7), the time-series standard deviation of alpha for the average firm is 2.09%. Thus, our conditional model picks up large fluctuations in pricing errors. Panel A further shows that for some firms variation in alpha is much more pronounced than for others. The first-order autocorrelation coefficients in panel B show that for most stocks pricing errors are quite persistent, irrespective of the specification used for betas. Results in panel C show that the betas produced by the model that only includes realized betas fluctuate more strongly over time than those obtained from the other two specifications. In the conditional beta model the standard deviation of beta is 0.30 for the average firm. The strong time-series variation in market risk exposure provides support for the use of conditional asset pricing models. However, the time-series variation in beta differs substantially across stocks, as the cross-sectional standard deviation equals 0.15.

As expected, panel D reveals that betas produced by the realized beta model are most persistent due to its autoregressive nature, while those generated by the fundamental beta model are least persistent. The conditional model leads to intermediate results, with an average first-order autocorrelation coefficient of 0.87. This indicates that our conditional betas represent the optimal trade-off between realized betas, which move too slow, and fundamentals-based betas, which change too fast. Although the firm-specific betas generated by our model are fairly persistent, Ang and Chen (2007) find that the monthly autocorrelation of the conditional

18

betas of value and growth portfolios equals 0.99 and Ghysels and Jacquier (2006) report that quarterly industry portfolio betas have strong autocorrelation on the order of 0.95. These findings suggest that portfolio betas are more persistent than individual stock betas. Finally, panel E displays important differences in the average correlation between alpha and beta across the three specifications. In particular, although this correlation equals -0.51 for the fundamental beta model, it is -0.06 for the model that only includes realized betas. The average correlation for the conditional model is -0.28. Hence, an increase in market risk tends to be related to a decrease in pricing errors. The large cross-sectional standard deviation (0.40) indicates considerable variation across firms in the relation between alpha and beta.

In sum, our empirical analysis reveals strong time variation in firm-specific conditional alphas and betas. Moreover, we find that realized betas and fundamentals-based betas exhibit different cross-sectional characteristics and time-series dynamics. Therefore, combining these two specifications captures different aspects of market beta dynamics and produces superior beta estimates. Furthermore, we document strong heterogeneity in risk dynamics across individual stocks. In the next section we examine the impact of these cross-sectional differences on the 25 size-B/M portfolios that are often used to test asset pricing models.

[Table 4 about here.]

4.5 Heterogeneity within Size-B/M Portfolios

An important assumption underlying the portfolio approach to testing asset pricing models is that the stocks in a particular portfolio share the same risk characteristics. In case of the widely used 25 portfolios sorted on firm size and book-to-market, it is assumed that firms are homogeneous in their exposure to market risk after controlling for size and B/M. However, a violation of this assumption can have serious asset pricing implications. In particular, if heterogeneity within size-B/M portfolios is large, portfolio averages conceal important information contained in the characteristics of individual stocks. In this section we therefore examine whether firms that are grouped together in a portfolio have similar risk and return characteristics.

We construct the 25 size-B/M portfolios according to the procedure outlined by Fama and French (1993). Specifically, at the end of June of each year, we sort all stocks independently into size and B/M quintiles. The 25 portfolios are then formed as the intersections of the size

and B/M quintiles. Subsequently, we calculate for every portfolio j and every month t the equally-weighted cross-sectional average of the posterior means of the firm-specific alphas, betas, and phis, and of the excess returns, market capitalizations, and B/M ratios of the stocks in that portfolio. Table 5 reports for each portfolio the time-series means of these cross-sectional averages. Consistent with prior studies (e.g., Fama and French (1996)), the small-growth portfolio has the lowest average return and a large, negative pricing error. In general, the average portfolio returns display a strong value premium but a weak size effect. Interestingly, sorting on firm size and B/M does not produce a wide spread in *average* market betas across portfolios, since all portfolio betas are close to one. However, Figure 9 shows that value and growth portfolios exhibit very different risk *dynamics*. In line with the findings of Ang and Chen (2007) and Franzoni (2007), the plot indicates that particularly during the 1980s value firms were less risky than growth stocks. These differences in risk dynamics also explain the sharp increase in the cross-sectional variation in betas during recessions, shown in Figure 5.[6]

[Figure 9 about here.]

The ϕ parameters in Table 5 show that realized betas are more important determinants of the conditional betas of large cap portfolios. In contrast, for small cap portfolios fundamentals-based betas have a stronger effect on conditional betas than realized betas. A possible explanation for this finding is that realized beta estimates for small cap portfolios are affected by non-synchronous trading effects, since small stocks are often less liquid.

[Table 5 about here.]

Table 6 displays the cross-sectional dispersion in the characteristics of the 25 size-B/M portfolios. For all characteristics we observe strong heterogeneity *within* portfolios. On average, the cross-sectional standard deviation of the alphas of the stocks included in a given portfolio is around 2%. Especially the firms that are grouped together in small-growth portfolios have significantly different pricing errors. Thus, even though for these portfolios average alphas are negative, many firms in the portfolios actually earn positive alphas. The cross-sectional standard deviation of the betas of the firms in a given portfolio is around 0.40. Clearly, the

---

[6]We also formed 10 portfolios based on momentum and find that winner and loser stocks also display opposite market risk dynamics. In general, winners have higher betas than losers.

assumption that stocks in the same portfolio share similar risk characteristics is violated. Table 6 also reports strong cross-sectional heterogeneity in ϕ within portfolios. This implies that for some firms in a portfolio conditional betas are mainly driven by realized betas whereas for others fundamentals-based betas are more important.

[Table 6 about here.]

5 Conclusion

In this paper we measure the market beta dynamics of individual stocks. We set up a Bayesian panel data model that allows a flexible specification of firm-specific betas. Our panel approach combines high-frequency return data and cross-sectional firm-level data to estimate betas with greater precision. By specifying hierarchical priors we can capture the cross-sectional heterogeneity in market betas without the need to estimate a large number of parameters. We find that modeling conditional betas as a linear combination of lagged fundamentals-based betas and past realized betas is preferred over traditional estimators of betas that are based on conditioning variables or data-driven filters alone. Specifically, because fundamentals-based betas and realized betas exhibit very different time-series dynamics and cross-sectional characteristics, combining both specifications captures different aspects of market beta. Moreover, in contrast to rolling window estimators of realized betas that give equal weight to past returns, we find that it is preferable to give more weight to recent observations.

We show that our model produces superior firm-specific beta forecasts, which we use to construct optimal portfolios subject to target beta constraints. Our empirical results further confirm the prediction of conditional asset pricing theory that individual stocks exhibit different risk dynamics. In particular, we find that the cross-sectional variation in beta increases sharply during recessions. Our findings also reveal strong cross-sectional heterogeneity in firm-specific alphas and betas *within* the 25 size-B/M sorted portfolios that are commonly used to test asset pricing models. Hence, the traditional portfolio approach to asset pricing, which ignores these differences within portfolios, leads to imprecise and flawed estimates of alphas and betas.

A Posterior Distributions

A.1 Joint Posterior Distribution

The joint posterior density is proportional to the product of the likelihood function and the prior distributions, $p(\theta|y) \propto p(y|\theta)p(\theta)$. Defining α_{it} and β_{it} as in section 1 and substituting the prior densities specified in section 2.1, the joint posterior is given by:

$$
\begin{aligned}
p(\theta|y) =&\, p(\alpha_{0i}, \sigma_{\alpha_0}^2, \alpha_1, \Omega_{\alpha_1}^{-1}, \phi_i, \sigma_\phi^2, \kappa_2, \delta_0, \sigma_{\delta_0}^2, \delta_1, \Omega_{\delta_1}^{-1}, \sigma_\eta^2, \sigma_{\epsilon_i}^2|y) \\
\propto&\, \prod_{i=1}^{N} \left(\sigma_{\epsilon_i}^2\right)^{-\frac{T_i}{2}} \exp\left[-\frac{1}{2\sigma_{\epsilon_i}^2}(r_{it} - \alpha_{it} - \beta_{it}r_{Mt})(r_{it} - \alpha_{it} - \beta_{it}r_M)\right] \\
\times&\, \prod_{i=1}^{N} (\sigma_{\alpha_0}^2)^{-\frac{1}{2}} \exp\left[-\frac{\alpha_{0i}^2}{2\sigma_{\alpha_0}^2}\right] \times \left(\sigma_{\alpha_0}^{-2}\right)^{A_{\alpha_0}+1} \exp\left[-\sigma_{\alpha_0}^{-2}B_{\alpha_0}\right] \\
\times&\, |\Omega_{\alpha_1}^{-1}|^{\frac{1}{2}} \exp\left[-\frac{1}{2}\alpha_1'\Omega_{\alpha_1}^{-1}\alpha_1\right] \times |\Omega_{\alpha_1}^{-1}|^{\frac{\psi_{\alpha_1}-(L+LM)-1}{2}} \exp\left[-\frac{1}{2}tr\left([\psi_{\alpha_1}S_{\alpha_1}]\Omega_{\alpha_1}^{-1}\right)\right] \\
\times&\, \prod_{i=1}^{N}(\sigma_\phi^2)^{-\frac{1}{2}} \exp\left[-\frac{(\phi_i - 0.5)^2}{2\sigma_\phi^2}\right] \times \left(\sigma_\phi^{-2}\right)^{A_\phi+1} \exp\left[-\sigma_\phi^{-2}B_\phi\right] \\
\times&\, (\sigma_{\delta_0}^2)^{-\frac{1}{2}} \exp\left[-\frac{\delta_0^2}{2\sigma_{\delta_0}^2}\right] \times \left(\sigma_{\delta_0}^{-2}\right)^{A_{\delta_0}+1} \exp\left[-\sigma_{\delta_0}^{-2}B_{\delta_0}\right] \\
\times&\, |\Omega_{\delta_1}^{-1}|^{\frac{1}{2}} \exp\left[-\frac{1}{2}\delta_1'\Omega_{\delta_1}^{-1}\delta_1\right] \times |\Omega_{\delta_1}^{-1}|^{\frac{\psi_{\delta_1}-(L+LM)-1}{2}} \exp\left[-\frac{1}{2}tr\left([\psi_{\delta_1}S_{\delta_1}]\Omega_{\delta_1}^{-1}\right)\right] \\
\times&\, \left(\sigma_\eta^{-2}\right)^{A_\eta+1} \exp\left[-\sigma_\eta^{-2}B_\eta\right] \times \prod_{i=1}^{N} \left(\sigma_{\epsilon_i}^{-2}\right)^{A_\epsilon+1} \exp\left[-\sigma_{\epsilon_i}^{-2}B_\epsilon\right].
\end{aligned}
$$

A.2 Conditional Posterior Distributions

In order to implement the Gibbs sampler we need to derive the full conditional posterior densities for each block of parameters. The conditional densities can be derived from the joint posterior density by ignoring all terms that do not depend on the parameters of interest and then treating the parameters considered to be known as constants. We then obtain the conditional density for the parameters of interest by rearranging the remaining terms into the kernel of a known distribution. We partition the parameter vector θ into the following blocks:

$\theta^{(1)}$: MIDAS weight parameter: (κ_2)

$\theta^{(2)}$: Individual effects: (α_{0i})

$\theta^{(3)}$: Fundamental alpha parameters: (α_1)

$\theta^{(4)}$: Fundamental beta parameters: (δ_0, δ_1)

$\theta^{(5)}$: Autoregressive beta parameter: (ϕ_i)

$\theta^{(6)}$: Variance and covariance parameters: $(\sigma_{\alpha_0}^2, \Omega_{\alpha_1}^{-1}, \sigma_\phi^2, \sigma_{\delta_0}^2, \Omega_{\delta_1}^{-1}, \sigma_\eta^2, \sigma_{\epsilon_i}^2)$

To generate samples from the conditional posterior of $\theta^{(1)}$ we use the Metropolis-Hastings algorithm. The conditional posteriors for all other blocks have convenient functional forms. Therefore, we use the Gibbs sampler to iteratively draw from the conditional densities of $\theta^{(2)}$, $\theta^{(3)}$, $\theta^{(4)}$, $\theta^{(5)}$, $\theta^{(6)}$. For notational convenience we rewrite the model in matrix form as

$$r_i = \alpha_{0i}\iota_{T_i} + ZBC_i\alpha_1 + r_M b_i \phi_i + r_M(\delta_0 \iota_T + ZBC_i\delta_1)(1 - \phi_i) + r_M \eta_i + \epsilon_i, \qquad (9)$$

where r_i is an $T \times 1$ vector of excess returns, ZBC_i is an $T \times (L + LM)$ matrix of conditioning variables, r_M an $T \times T$ diagonal matrix of excess market returns, and b_i an $T \times 1$ vector of length T of realized betas. η_i and ϵ_i are $T \times 1$ vectors of zero-mean, normally distributed idiosyncratic shocks to betas and to returns, respectively.

Since the δ_0 and δ_1 parameters are both in block $\theta^{(4)}$ and have independent priors, we can rewrite $\delta_0 \iota_T + ZBC_i\delta_1$ as $W_i\delta$, where W_i is the $T \times (1 + L + LM)$ matrix of the constant term and the conditioning variables. We combine the corresponding precisions $\sigma_{\delta_0}^{-2}$ and $\Omega_{\delta_1}^{-1}$ into the matrix Ω_δ^{-1}.

A.2.1 Metropolis-Hastings algorithm to draw κ_2

Since we implement a change of variable, $\kappa_2 = 1 + 25\tilde{\kappa}_2$, we need to draw values for $\tilde{\kappa}_2$. Because the conditional posterior density for $\tilde{\kappa}_2$ does not take a standard form, we cannot use the Gibbs sampler to draw values for $\tilde{\kappa}_2$. Instead, we employ the Metropolis-Hastings algorithm, which is a general accept-reject algorithm. In fact, Gelman, Carlin, Stern, and Rubin (2004) show that the Gibbs sampler is a special case of Metropolis-Hastings in which proposed parameter values are accepted with probability one. The M-H algorithm proceeds as follows.

First, a candidate value $\tilde{\kappa}_2^*$ is drawn from a proposal density $q(\tilde{\kappa}_2)$. We apply the Independence Chain M-H algorithm, in which the proposal density is independent across draws. We choose a Beta(1,3) proposal density, which has a mean of 0.25 and standard deviation equal to 0.19. Because the proposal density is not identical to the posterior density, the M-H algorithm does not accept all proposal draws. When a proposal is rejected the parameter value is set equal to the current value. Draws are accepted according to the following probability

$$\pi(\tilde{\kappa}_2^{(g-1)}, \tilde{\kappa}_2^*) = min\left\{1, \frac{p(\tilde{\kappa}_2^*|y)q(\tilde{\kappa}_2^{(g-1)})}{p(\tilde{\kappa}_2^{(g-1)}|y)q(\tilde{\kappa}_2^*)}\right\}. \tag{10}$$

This approach ensures that candidate draws with a high posterior density have a higher probability of being accepted than draws with a low posterior density. Repeating this procedure produces the required sequence of draws from the posterior distribution.

A.2.2 Conditional posterior α_{0i}

Using Bayes' theorem, we can write:

$$p\left(\alpha_{0i}|y\right) \propto p\left(y|\alpha_{0i}\right)p\left(\alpha_{0i}\right)$$

$$\propto \exp\left[-\frac{1}{2}Q^*\right],$$

$$\text{where}\quad Q^* = (r_i - \alpha_i - r_M\beta_i)'\,\Omega_{\epsilon_i}^{-1}(r_i - \alpha_i - r_M\beta_i) + \frac{\alpha_{0i}^2}{\sigma_{\alpha_0}^2}$$

$$= (X_{\alpha_{0i}} - \alpha_{0i}\iota_{T_i})'\,\Omega_{\epsilon_i}^{-1}\left(X_{\alpha_{0i}} - \alpha_{0i}\iota_{T_i}\right) + \frac{\alpha_{0i}^2}{\sigma_{\alpha_0}^2}$$

$$= Q_1^* + Q_2^*,$$

$$\text{with}\quad Q_1^* = \frac{(\alpha_{0i} - \bar{\alpha}_{0i})^2}{\sigma_{\alpha_{0i}}^2},$$

$$\text{and}\quad Q_2^* = X_{\alpha_{0i}}'\,\Omega_{\epsilon_i}^{-1}X_{\alpha_{0i}} - \frac{\bar{\alpha}_{0i}^2}{\sigma_{\alpha_{0i}}^2},$$

$$\text{and where}\quad X_{\alpha_{0i}} = r_i - \left(Z_i\alpha_1 + r_M b_i\phi_i + r_M W_i\delta(1 - \phi_i) + r_M\eta_i\right).$$

In the derivation of $p\left(\alpha_{0i}|\theta^{-(\alpha_{0i})}, y\right)$ all parameters in Q_2^* are known, so we can treat Q_2^* as a constant. Thus, $p\left(\alpha_{0i}|\theta^{-(\alpha_{0i})}, y\right)$ is proportional to $\exp[-\frac{1}{2}Q_1^*]$, which is the kernel of a normal density. Therefore,

$$\alpha_{0i}|\theta^{-(\alpha_{0i})}, y \sim N\left(\bar{\alpha}_{0i}, \sigma_{\alpha_{0i}}^2\right),$$

$$\text{with}\quad \bar{\alpha}_{0i} = \left[\iota_{T_i}'\,\Omega_{\epsilon_i}^{-1}\iota_{T_i} + \frac{1}{\sigma_{\alpha_0}^2}\right]^{-1}\left[\iota_{T_i}'\,\Omega_{\epsilon_i}^{-1}X_{\alpha_{0i}}\right],$$

$$\text{and}\quad \sigma_{\alpha_{0i}}^2 = \left[\iota_{T_i}'\,\Omega_{\epsilon_i}^{-1}\iota_{T_i} + \frac{1}{\sigma_{\alpha_0}^2}\right]^{-1}.$$

We derive all other conditional posteriors in the same way, which leads to the following results:

A.2.3 Conditional posterior α_1

$$\alpha_1 | \theta^{-(\alpha_1)}, y \sim N\left(\bar{\alpha}_1, V_{\alpha_1}\right),$$

$$\text{with} \quad \bar{\alpha}_1 = \left[\sum_{i=1}^{N} Z_i' \Omega_{\epsilon_i}^{-1} Z_i + \Omega_{\alpha_1}^{-1}\right]^{-1} \left[\sum_{i=1}^{N} Z_i' \Omega_{\epsilon_i}^{-1} X_{\alpha_{1i}}\right],$$

$$\text{and} \quad V_{\alpha_1} = \left[\sum_{i=1}^{N} Z_i' \Omega_{\epsilon_i}^{-1} Z_i + \Omega_{\alpha_1}^{-1}\right]^{-1},$$

$$\text{and where} \quad X_{\alpha_{1i}} = r_i - (\alpha_{0i} \iota_{T_i} + r_M b_i \phi_i + r_M W_i \delta (1 - \phi_i) + r_M \eta_i).$$

A.2.4 Conditional posterior δ

$$\delta | \theta^{-(\delta)}, y \sim N\left(\bar{\delta}, V_\delta\right),$$

$$\text{with} \quad \bar{\delta} = \left[\sum_{i=1}^{N} (1 - \phi_i) W_i' r_M' \Omega_{\epsilon_i}^{-1} r_M W_i (1 - \phi_i) + \Omega_\delta^{-1}\right]^{-1} \left[\sum_{i=1}^{N} (1 - \phi_i) W_i' r_M' \Omega_{\epsilon_i}^{-1} X_{\delta_i}\right],$$

$$\text{and} \quad V_\delta = \left[\sum_{i=1}^{N} (1 - \phi_i) W_i' r_M' \Omega_{\epsilon_i}^{-1} r_M W_i (1 - \phi_i) + \Omega_\delta^{-1}\right]^{-1},$$

and where $\quad X_{\delta_i} = r_i - (\alpha_{0i} \iota_{T_i} + Z_i \alpha_1 + r_M b_i \phi_i + r_M \eta_i).$

A.2.5 Conditional posterior ϕ_i

$$\phi_i | \theta^{-(\phi_i)}, y \sim N\left(\bar{\phi}_i, \sigma_{\phi_i}^2\right),$$

$$\text{with} \quad \bar{\phi}_i = \left[(r_M b_i - r_M W_i \delta)' \Omega_{\epsilon_i}^{-1} (r_M b_i - r_M W_i \delta) + \frac{1}{\sigma_\phi^2}\right]^{-1} \left[(r_M b_i - r_M W_i \delta)' \Omega_{\epsilon_i}^{-1} X_{\phi_i} + \frac{0.5}{\sigma_\phi^2}\right],$$

$$\text{and} \quad \sigma_{\phi_i}^2 = \left[(r_M b_i - r_M W_i \delta)' \Omega_{\epsilon_i}^{-1} (r_M b_i - r_M W_i \delta) + \frac{1}{\sigma_\phi^2}\right]^{-1},$$

and where $\quad X_{\phi_i} = r_i - (\alpha_{0i} \iota_{T_i} + Z_i \alpha_1 + r_M W_i \delta + r_M \eta_i).$

A.2.6 Conditional posteriors $\sigma_{\alpha_0}^2, \Omega_{\alpha_1}^{-1}, \Omega_{\delta}^{-1}, \sigma_{\phi}^2, \sigma_{\eta}^2, \sigma_{\epsilon_i}^2$

$$\sigma_{\alpha_0}^2 | \theta^{-(6)}, y \sim IG\left(\frac{N + 2A_{\alpha_0}}{2}, \frac{\sum_{i=1}^{N} \alpha_{0i}^2 + 2B_{\alpha_0}}{2}\right),$$

$$\Omega_{\alpha_1}^{-1} | \theta^{-(6)}, y \sim Wish\left(\left[\alpha_1\alpha_1' + (\psi_{\alpha_1}S_{\alpha_1})\right]^{-1}, \psi_{\alpha_1} + 1\right),$$

$$\Omega_{\delta}^{-1} | \theta^{-(6)}, y \sim Wish\left(\left[\delta\delta' + (\psi_{\delta}S_{\delta})\right]^{-1}, \psi_{\delta} + 1\right),$$

$$\sigma_{\phi}^2 | \theta^{-(6)}, y \sim IG\left(\frac{N + 2A_{\phi}}{2}, \frac{\sum_{i=1}^{N} \phi_i^2 + 2B_{\phi}}{2}\right),$$

$$\sigma_{\eta}^2 | \theta^{-(6)}, y \sim IG\left(\frac{\sum_{i=1}^{N} T_i + 2A_{\eta}}{2}, \frac{\sum_{i=1}^{N} \eta_i'\eta_i + 2B_{\eta}}{2}\right),$$

$$\sigma_{\epsilon_i}^2 | \theta^{-(6)}, y \sim IG\left(\frac{T_i + 2A_{\epsilon}}{2}, \frac{(r_i - \alpha_i - r_M\beta_i)'(r_i - \alpha_i - r_M\beta_i) + 2B_{\epsilon}}{2}\right).$$

References

Ang, A., and J. Chen, 2007, CAPM over the long run: 1926-2001, *Journal of Empirical Finance* 14, 1–40.

Avramov, D., and T. Chordia, 2006, Asset pricing models and financial market anomalies, *Review of Financial Studies* 19, 1001–1040.

Bauwens, L., M. Lubrano, and J. Richard, 1999, *Bayesian Inference in Dynamic Econometric Models* (Oxford University Press: Oxford).

Bollerslev, T., and B. Zhang, 2003, Measuring and modeling systematic risk in factor pricing models using high-frequency data, *Journal of Empirical Finance* 10, 533–558.

Brennan, M., T. Chordia, and A. Subrahmanyam, 1998, Alternative factor specifications, security characteristics, and the cross-section of expected stock returns, *Journal of Financial Economics* 49, 345–373.

Campbell, J., A. Lo, and A. MacKinlay, 1997, *The Econometrics of Financial Markets* (Princeton University Press: Princeton).

Chib, S., and E. Greenberg, 1995, Understanding the Metropolis-Hastings algorithm, *The American Statistician* 49, 327–335.

Christoffersen, P., K. Jacobs, and G. Vainberg, 2007, Forward-looking betas, Working Paper, McGill University.

Fama, E., and K. French, 1989, Business conditions and expected returns on stocks and bonds, *Journal of Financial Economics* 25, 23–49.

Fama, E., and K. French, 1993, Common risk factors in the returns on stocks and bonds, *Journal of Financial Economics* 33, 3–56.

Fama, E., and K. French, 1996, Multifactor explanations of asset pricing anomalies, *Journal of Finance* 51, 55–84.

Fama, E., and K. French, 1997, Industry costs of equity, *Journal of Financial Economics* 43, 153–193.

Fama, E., and J. MacBeth, 1973, Risk, return and equilibrium: Empirical tests, *Journal of Political Economy* 71, 607–636.

Ferson, W., and C. Harvey, 1999, Conditioning variables and the cross section of stock returns, *Journal of Finance* 54, 1325–1360.

Franzoni, F., 2007, The changing nature of systematic risk, Working Paper, HEC School of Management.

Gelman, A., J. Carlin, H. Stern, and D. Rubin, 2004, *Bayesian Data Analysis* (Chapman and Hall/CRC: Boca Raton, Florida) 2nd edn.

Geweke, J., 2005, *Contemporary Bayesian Econometrics and Statistics* (Wiley: New York).

Ghysels, E., 1998, On stable factor structures in the pricing of risk: Do time-varying betas help or hurt?, *Journal of Finance* 53, 549–573.

Ghysels, E., and E. Jacquier, 2006, Market beta dynamics and portfolio efficiency, Working paper, University of North Carolina and HEC Montreal.

Ghysels, E., P. Santa-Clara, and R. Valkanov, 2005, There is a risk-return trade-off after all, *Journal of Financial Economics* 76, 509–548.

Gomes, J., L. Kogan, and L. Zhang, 2003, Equilibrium cross-section of returns, *Journal of Political Economy* 111, 693–732.

Hoberg, G., and I. Welch, 2007, Aged and recent market betas in securities prices, Working Paper, University of Maryland and Brown University.

Hsiao, C., and M. Pesaran, 2007, Random coefficient panel data models, in L. Matyas, and P. Sevestre, ed.: *The Econometrics of Panel Data* (Kluwer Academic Publishers) 3rd edn.

Jagannathan, R., and Z. Wang, 1996, The conditional CAPM and cross-section of expected returns, *Journal of Finance* 51, 3–53.

Jostova, G., and A. Philipov, 2005, Bayesian analysis of stochastic betas, *Journal of Financial and Quantitative Analysis* 40.

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, 1238–1287.

Lewellen, J., 1999, The time-series relations among expected return, risk, and book-to-market, *Journal of Financial Economics* 54, 5–43.

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

Petkova, R., and L. Zhang, 2005, Is value riskier than growth?, *Journal of Financial Economics* 78, 187–202.

Santos, T., and P. Veronesi, 2004, Conditional betas, Working paper, Columbia University and University of Chicago.

Shanken, J., 1990, Intertemporal asset pricing: An empirical investigation, *Journal of Econometrics* 45, 99–120.

Zhang, L., 2005, The value premium, *Journal of Finance* 60, 67–103.

Table 1: **Summary Statistics for Firm Characteristics and Macroeconomic Variables**

This table presents descriptive statistics for stock returns, firm characteristics, and macroeconomic variables for 510 months from July 1964 through December 2006. Panel A reports the mean, median, standard deviation and 5th, 25th, 75th, and 95th percentile values of firm characteristics for a total of 5,017 stocks over the full sample period and an average of 1,815 stocks per month. We include a stock in the sample for a given month t if it satisfies the following criteria. First, its return in the current month, t, and over the past 36 months has to be available. Second, data should be available in month $t-1$ for size as measured by market capitalization, for the book-to-market ratio, and for turnover. We exclude firms with negative book-to-market equity. XRET is the return in excess of the risk-free rate, MV represents the market capitalization in billions of dollars, and BM is the book-to-market ratio, for which values smaller than the 0.5th percentile and values greater than the 99.5th percentile are set equal to the 0.5th percentile and 99.5th percentile values, respectively. MOM is the cumulative return over the twelve months prior to the current month. TURN is monthly share turnover, defined as trading volume divided by the numbers of shares outstanding. Panel B shows the the mean, median, standard deviation and 5th, 25th, 75th, and 95th percentile values of macroeconomic variables. DEF is the default spread, defined as the yield differential between bonds rated BAA by Moody's and bonds with a Moody's rating of AAA. DY is the dividend yield on the value-weighted CRSP index. The dividend yield is calculated as the sum of the dividends paid on the index in the previous year divided by the current level of the index. TBILL is the one-month Treasury bill rate. TERM is the term spread, defined as the yield difference between ten-year and one-year Treasury bonds.

	Mean	Std. dev.	5th	25th	Median	75th	95th
Panel A: Firm characteristics							
XRET (%)	0.69	12.31	-17.48	-5.99	-0.16	6.42	21.29
MV ($ billions)	1.59	5.60	0.01	0.03	0.16	0.81	6.67
BM	0.96	0.82	0.18	0.44	0.75	1.22	2.45
MOM (%)	14.65	49.29	-49.12	-14.24	8.60	34.46	96.30
TURN (%)	5.19	6.36	0.47	1.41	3.06	6.37	17.15
Panel B: Macroeconomic variables							
DEF (%)	1.02	0.43	0.55	0.73	0.90	1.21	1.92
DY (%)	3.01	1.10	1.30	2.02	2.96	3.77	4.84
TBILL (%)	5.69	2.70	1.56	4.08	5.16	6.96	10.57
TERM (%)	0.85	1.14	-1.14	0.08	0.78	1.69	2.83

Table 2: **Posterior Distribution of Fundamentals-Based Beta Parameters**

This table reports the Bayesian posterior distribution of the determinants of the fundamentals-based beta, which is parameterized as a linear function of lagged firm characteristics and business cycle variables,

$$\beta_{it-1}^* = \delta_0 + \delta_1'[Z_{it-1} \otimes BC_{t-1}],$$

where Z_{it-1} is a vector that contains L firm characteristics and BC_{t-1} is a vector that contains a constant and M business cycle variables. The table presents the mean, median, standard deviation and 5th, 25th, 75th, and 95th percentile values of the posterior distribution, based on 5,000 iterations of the Gibbs sampler and a burn-in period of 1,000 iterations.

	Mean	Std. dev.	5th	25th	Median	75th	95th
Intercept	1.15	0.02	1.12	1.14	1.15	1.17	1.19
MV	-0.12	0.02	-0.15	-0.13	-0.12	-0.11	-0.09
BM	0.85	0.07	0.72	0.79	0.84	0.90	0.97
MOM	0.18	0.15	-0.07	0.08	0.18	0.28	0.42
TURN	-0.02	0.01	-0.03	-0.02	-0.02	-0.01	-0.01
MV*TBILL	0.01	0.01	0.01	0.01	0.01	0.02	0.02
MV*TERM	0.03	0.01	0.02	0.02	0.03	0.03	0.03
MV*DEF	0.00	0.01	-0.02	-0.01	0.00	0.00	0.01
MV*DY	0.00	0.00	0.00	0.00	0.00	0.01	0.01
BM*TBILL	-0.06	0.01	-0.08	-0.07	-0.06	-0.05	-0.03
BM*TERM	-0.08	0.02	-0.11	-0.10	-0.08	-0.07	-0.05
BM*DEF	-0.50	0.05	-0.59	-0.54	-0.50	-0.47	-0.42
BM*DY	-0.04	0.01	-0.05	-0.04	-0.04	-0.03	-0.02
MOM*TBILL	0.01	0.03	-0.04	-0.01	0.01	0.03	0.06
MOM*TERM	-0.06	0.05	-0.14	-0.09	-0.06	-0.03	0.02
MOM*DEF	-0.16	0.07	-0.27	-0.20	-0.15	-0.11	-0.04
MOM*DY	0.02	0.03	-0.03	0.00	0.02	0.04	0.07
TURN*TBILL	0.00	0.00	0.00	0.00	0.00	0.00	0.00
TURN*TERM	-0.01	0.00	-0.01	-0.01	-0.01	-0.01	-0.01
TURN*DEF	0.06	0.01	0.05	0.06	0.06	0.07	0.08
TURN*DY	0.00	0.00	-0.01	0.00	0.00	0.00	0.00

Table 3: **Forecasting Power of Alternative Approaches to Modeling Betas**

This table reports regression estimates for optimal portfolios that minimize the sum of squared portfolio weights subject to the constraint that the one-month ahead forecast of the portfolio beta equals the target beta. The portfolio are constructed at the end of month $t-1$ using the forecasts of the stock-specific betas at time t. We compare the forecasting power of four different beta specifications, i.e., the full conditional beta given by equation (2), the realized beta in equation (3), the fundamentals-based beta in equation (5), and the static beta in the unconditional CAPM. We measure the accuracy of the beta forecasts by regressing the time-series of portfolio returns on the market index and on the interaction terms between the market index and the forecasts of the portfolio beta produced by the beta specifications other than the one used to construct the portfolio,

$$r_{pt} = \alpha_p + \beta_p r_{Mt} + \gamma_p' X_{pt-1} r_{Mt} + \epsilon_{pt}.$$

Test 1 investigates whether the realized portfolio beta equals the target beta. Test 2 examines whether deviations from the target beta are predictable by the other beta specifications. Panel A shows the regression results for the portfolios with a target beta of one and panel B shows the estimates for portfolios with a target beta of zero. t-statistics are in parentheses and computed using Newey-West heteroskedasticity and autocorrelation robust standard errors. The t-statistic for the coefficient on r_M in panel A corresponds to the null hypothesis that this coefficient equals one. The sample period is September 1964 through December 2006 (508 months).

	Conditional		Fundamental		Realized		Static	
	Test 1	Test 2	Test 1	Test 2	Test 1	Test 2	Test 1	Test 2
Panel A: Target Portfolio Beta = 1								
Intercept	0.13	0.14	0.14	0.15	0.13	0.14	0.23	0.24
	(1.16)	(1.28)	(1.11)	(1.31)	(1.15)	(1.28)	(1.75)	(1.97)
r_M	1.00	1.01	0.99	1.02	0.99	0.97	1.06	1.08
	(0.10)	(0.25)	(0.17)	(0.66)	(0.43)	(1.02)	(1.50)	(2.28)
$\beta_{conditional} * r_M$				68.16		32.73		84.74
				(6.55)		(2.24)		(4.94)
$\beta_{fundamental} * r_M$		-7.55				13.05		-3.91
		(-1.64)				(2.02)		(-2.35)
$\beta_{realized} * r_M$		5.62		40.16				46.97
		(1.13)		(6.62)				(5.06)
$\beta_{static} * r_M$		-0.88		0.53		3.58		
		(-1.26)		(1.60)		(1.88)		
Adj. R^2(%)	81.5	81.8	78.3	81.4	80.6	81.0	77.7	79.9
Panel B: Target Portfolio Beta = 0								
Intercept	0.07	0.06	0.07	0.05	0.06	0.07	0.01	0.01
	(3.92)	(4.48)	(3.05)	(3.46)	(3.25)	(3.97)	(0.85)	(2.26)
r_M	0.02	0.00	0.05	0.00	0.07	0.00	0.03	0.01
	(3.42)	(0.47)	(5.19)	(0.41)	(9.80)	(-0.20)	(6.41)	(2.64)
$\beta_{conditional} * r_M$				2.86		3.24		1.20
				(3.02)		(4.65)		(4.00)
$\beta_{fundamental} * r_M$		-0.80				-1.16		-0.42
		(-3.23)				(-4.33)		(-2.86)
$\beta_{realized} * r_M$		-0.50		0.26				-0.01
		(-1.36)		(0.48)				(-0.06)
$\beta_{static} * r_M$		0.28		-0.76		0.11		
		(1.12)		(-2.35)		(0.53)		
Adj. R^2(%)	8.3	23.7	21.7	50.4	44.8	54.0	32.1	69.7

Table 4: **Time-Series Dynamics of Alphas and Betas**

This table reports time-series characteristics of conditional alphas and betas produced by three different specification for beta. Specifically, we compare the conditional beta given by equation (2), the realized beta in equation (3), and the fundamentals-based beta in equation (5). We first calculate at each iteration of the Gibbs sampler the conditional alpha and beta for firm i at time t. Subsequently, we compute for every firm and every date the average alpha and beta over all iterations (posterior mean). We then compute the time-series standard deviation and first-order autocorrelation of alpha and beta and the correlation between alpha and beta for every firm i. Panels A and B report the mean, median, standard deviation and 5th, 25th, 75th, and 95th percentile values of the cross-sectional distribution of the time-series standard deviation and autocorrelation of alphas, respectively. Panels C and D present the cross-sectional distribution of the time-series standard deviation and autocorrelation of betas, respectively. Panel E shows the cross-sectional distribution of the correlation between alphas and betas.

Model	Mean	Std. dev.	5th	25th	Median	75th	95th
Panel A: Time-Series Variation in Alphas							
Conditional Beta	2.10	1.41	0.46	1.02	1.81	2.86	4.71
Realized Beta	1.78	1.17	0.40	0.89	1.54	2.40	3.93
Fundamental Beta	2.45	1.65	0.52	1.18	2.11	3.33	5.50
Panel B: Autocorrelation in Alphas							
Conditional Beta	0.82	0.15	0.52	0.77	0.86	0.92	0.96
Realized Beta	0.82	0.15	0.52	0.77	0.86	0.92	0.96
Fundamental Beta	0.81	0.15	0.51	0.77	0.86	0.92	0.96
Panel C: Time-Series Variation in Betas							
Conditional Beta	0.30	0.15	0.11	0.19	0.27	0.37	0.58
Realized Beta	0.41	0.20	0.15	0.27	0.37	0.51	0.78
Fundamental Beta	0.30	0.16	0.11	0.19	0.27	0.38	0.59
Panel D: Autocorrelation in Betas							
Conditional Beta	0.87	0.13	0.59	0.84	0.92	0.95	0.98
Realized Beta	0.93	0.06	0.81	0.92	0.95	0.97	0.98
Fundamental Beta	0.70	0.19	0.32	0.60	0.74	0.84	0.92
Panel E: Correlation between Alphas and Betas							
Conditional Beta	-0.28	0.40	-0.84	-0.61	-0.34	0.01	0.46
Realized Beta	-0.06	0.34	-0.60	-0.30	-0.06	0.18	0.52
Fundamental Beta	-0.51	0.39	-0.92	-0.79	-0.63	-0.34	0.35

Table 5: **Average Characteristics of 25 Size-B/M Portfolios**

This table presents the time-series averages of characteristics of 25 size-B/M sorted portfolios. The portfolios are constructed annually by sorting stocks independently into size and B/M quintiles at the end of June. The 25 portfolios are then formed as the intersections of these five size and B/M quintiles. Subsequently, we calculate for every portfolio j at every time t the equally weighted cross-sectional average of the posterior means of the firm-specific alphas, betas, and phis, and of the excess returns, market capitalizations, and B/M ratios of the stocks in the portfolio. Firm size and B/M are expressed as natural logarithms and as deviations from their cross-sectional means. The table shows for every portfolio the time-series means of these cross-sectional averages.

Size	Book-to-Market equity (B/M) quintiles									
quintiles	Low	2	3	4	High	Low	2	3	4	High
			Alpha					Beta		
Small	-2.28	-0.40	0.17	0.65	1.28	1.23	1.08	0.99	0.90	0.84
2	-1.27	-0.06	0.39	0.86	1.42	1.31	1.13	1.03	0.97	0.95
3	-0.94	0.02	0.40	0.81	1.38	1.32	1.13	1.02	0.97	1.01
4	-0.96	-0.06	0.36	0.74	1.41	1.24	1.10	1.00	0.94	0.97
Big	-0.88	-0.20	0.17	0.58	1.17	1.19	1.07	0.99	0.93	0.96
			Phi					Return		
Small	0.39	0.41	0.42	0.45	0.46	0.18	0.54	0.67	0.97	0.93
2	0.47	0.49	0.51	0.52	0.52	0.72	0.83	0.80	0.83	0.85
3	0.55	0.54	0.55	0.58	0.61	0.67	0.69	0.76	0.95	0.87
4	0.59	0.57	0.60	0.63	0.64	0.65	0.62	0.72	0.84	0.94
Big	0.61	0.61	0.65	0.70	0.72	0.48	0.46	0.59	0.72	0.84
			Size					B/M		
Small	-2.58	-2.57	-2.58	-2.61	-2.81	-1.12	-0.30	0.07	0.40	1.01
2	-1.09	-1.08	-1.08	-1.10	-1.17	-1.06	-0.30	0.07	0.39	0.90
3	0.03	0.01	-0.02	-0.02	-0.05	-1.01	-0.30	0.06	0.38	0.88
4	1.13	1.12	1.10	1.08	1.04	-1.01	-0.31	0.06	0.38	0.83
Big	2.92	2.69	2.64	2.55	2.46	-1.06	-0.31	0.05	0.37	0.80

Table 6: **Cross-Sectional Variation in Characteristics of 25 Size-B/M Portfolios**

This table presents the time-series averages of the cross-sectional variation in characteristics of 25 size-B/M sorted portfolios. The portfolios are constructed annually by sorting stocks independently into size and B/M quintiles at the end of June. The 25 portfolios are then formed as the intersections of these five size and B/M quintiles. Subsequently, we calculate for every portfolio j at every time t the cross-sectional standard deviation of the posterior means of the firm-specific alphas, betas, and phis, and of the excess returns, market capitalizations, and B/M ratios of the stocks in the portfolio. Firm size and B/M are expressed as natural logarithms and as deviations from their cross-sectional means. The table shows for every portfolio the time-series means of these cross-sectional standard deviations.

Size	Book-to-Market equity (B/M) quintiles									
quintiles	Low	2	3	4	High	Low	2	3	4	High
			Alpha					Beta		
Small	3.67	2.57	2.38	2.24	2.51	0.47	0.39	0.38	0.38	0.40
2	3.24	2.35	2.11	2.01	2.33	0.45	0.39	0.38	0.37	0.38
3	2.96	2.19	1.93	1.88	2.14	0.46	0.40	0.39	0.38	0.38
4	2.57	1.91	1.73	1.58	1.74	0.43	0.38	0.38	0.37	0.37
Big	2.17	1.66	1.47	1.35	1.51	0.35	0.31	0.32	0.32	0.30
			Phi					Return		
Small	0.32	0.32	0.32	0.32	0.31	15.56	13.84	13.11	12.50	13.83
2	0.33	0.32	0.31	0.30	0.31	12.98	11.16	10.40	10.13	11.34
3	0.32	0.32	0.31	0.30	0.28	11.03	9.83	9.04	8.89	9.82
4	0.31	0.30	0.29	0.30	0.28	9.26	8.26	7.83	7.51	8.33
Big	0.26	0.29	0.29	0.28	0.24	7.27	6.93	6.57	6.35	6.56
			Size					B/M		
Small	0.61	0.60	0.62	0.63	0.74	0.59	0.12	0.09	0.10	0.35
2	0.34	0.34	0.34	0.35	0.35	0.50	0.12	0.09	0.10	0.30
3	0.30	0.30	0.30	0.31	0.31	0.45	0.12	0.09	0.10	0.29
4	0.34	0.34	0.33	0.33	0.33	0.43	0.13	0.09	0.10	0.24
Big	0.88	0.74	0.72	0.64	0.55	0.42	0.13	0.09	0.10	0.20

Figure 1: **Cross-Sectional Distribution of Posterior Mean of Phi**

This figure shows the cross-sectional distribution of the posterior mean of the parameter ϕ_i, which measures the proportion of beta explained by past realized beta,

$$\beta_{it} = \phi_i b_{it-1} + (1 - \phi_i)\beta_{it-1}^* + \eta_{it},$$

where b_{it-1} is firm's i past realized beta and where β_{it-1}^* is the lagged fundamentals-based beta of firm i. We give more weight to the ϕ_i that are estimated with higher precision by weighting each posterior mean by its posterior precision. This figure shows the cross-sectional distribution of these precision-weighted posterior means.
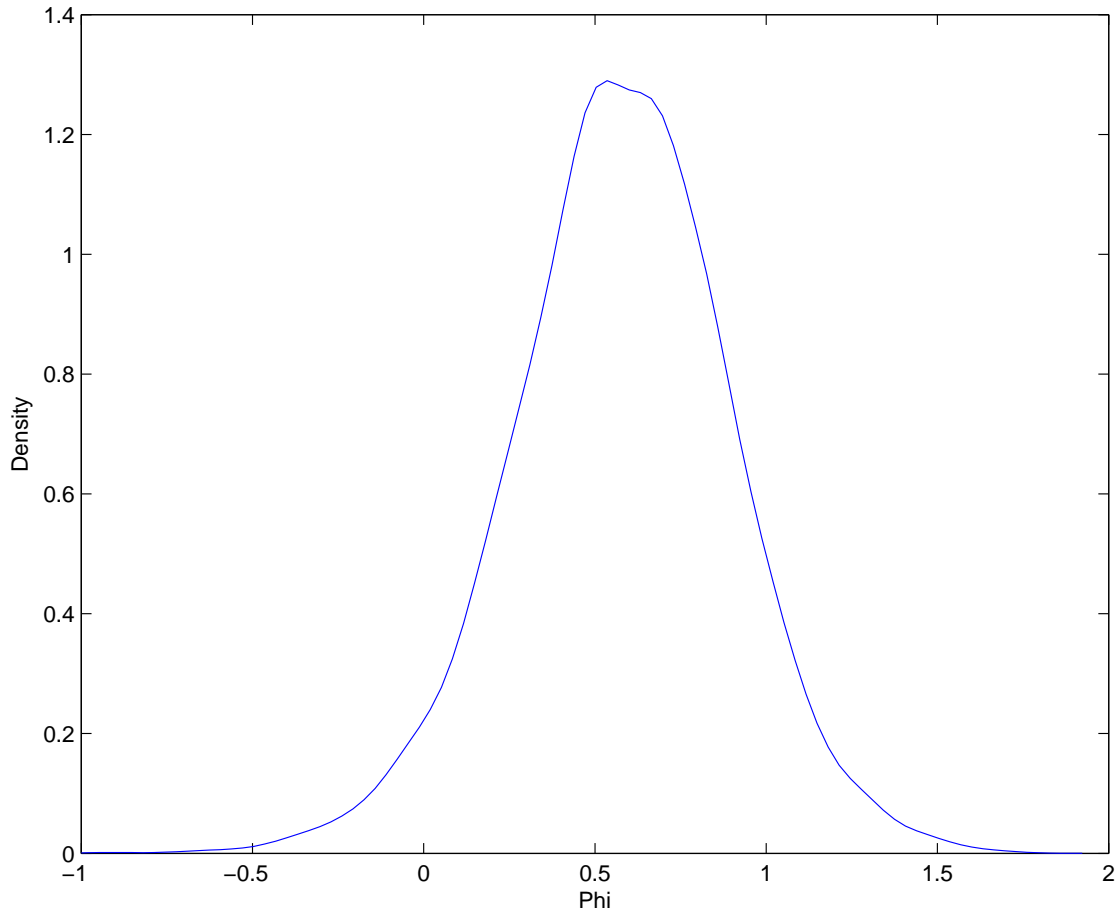
Figure 2: **Optimal versus Equal Weighting Scheme for Estimating Beta**

This figure compares the equal weights in the traditional rolling window estimator of realized betas to the weights implied by the estimated MIDAS weighting function in equation (4) for the model in equation (7). We set the maximum window length equal to 250 trading days.
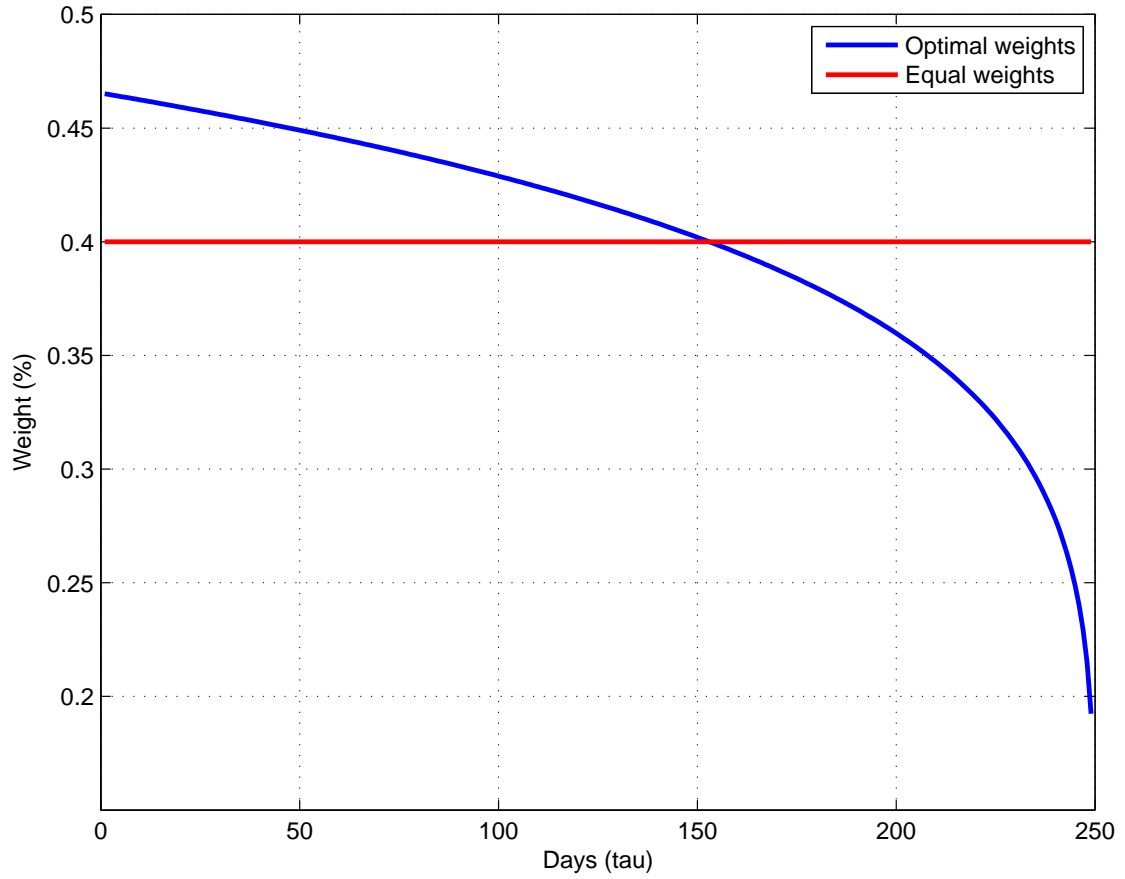
Figure 3: **Cross-Sectional Distribution of Firm Betas**

This figure shows the cross-sectional distribution of average firm betas. We first calculate at each iteration of the Gibbs sampler the beta for firm i at time t based on the model in equation (7). Subsequently, we compute the time-series averages of these conditional betas. We then calculate for each firm the posterior mean of its time-series average beta. We give more weight to betas that are estimated with higher precision by weighting each posterior mean by its posterior precision. This figure shows the cross-sectional distribution of these precision-weighted posterior means.
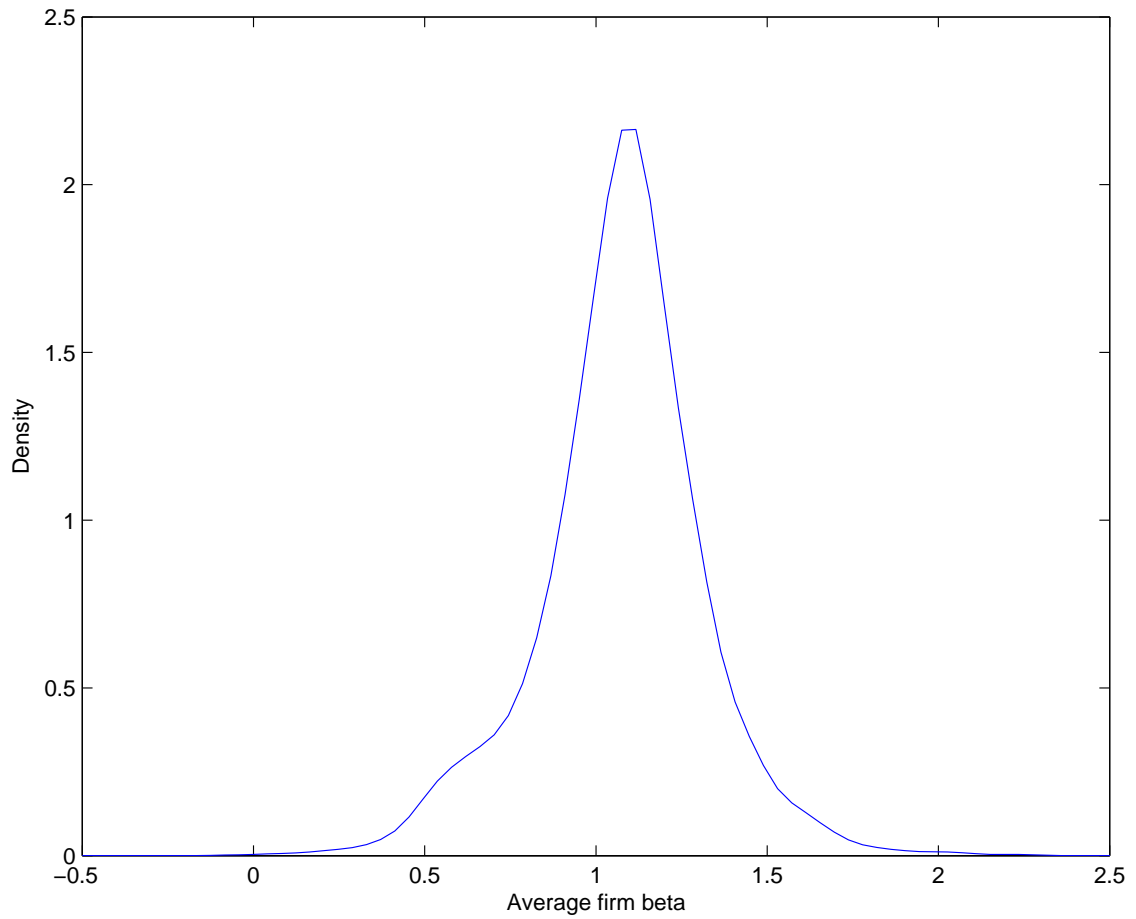
Figure 4: **Cross-Sectional Distribution of Firm Alphas**

This figure shows the cross-sectional distribution of average firm alphas. We first calculate at each iteration of the Gibbs sampler the alpha for firm i at time t based on the model in equation (7). Subsequently, we compute the time-series averages of these conditional alphas. We then calculate for each firm the posterior mean of its time-series average alpha. We give more weight to alphas that are estimated with higher precision by weighting each posterior mean by its posterior precision. This figure shows the cross-sectional distribution of these precision-weighted posterior means.
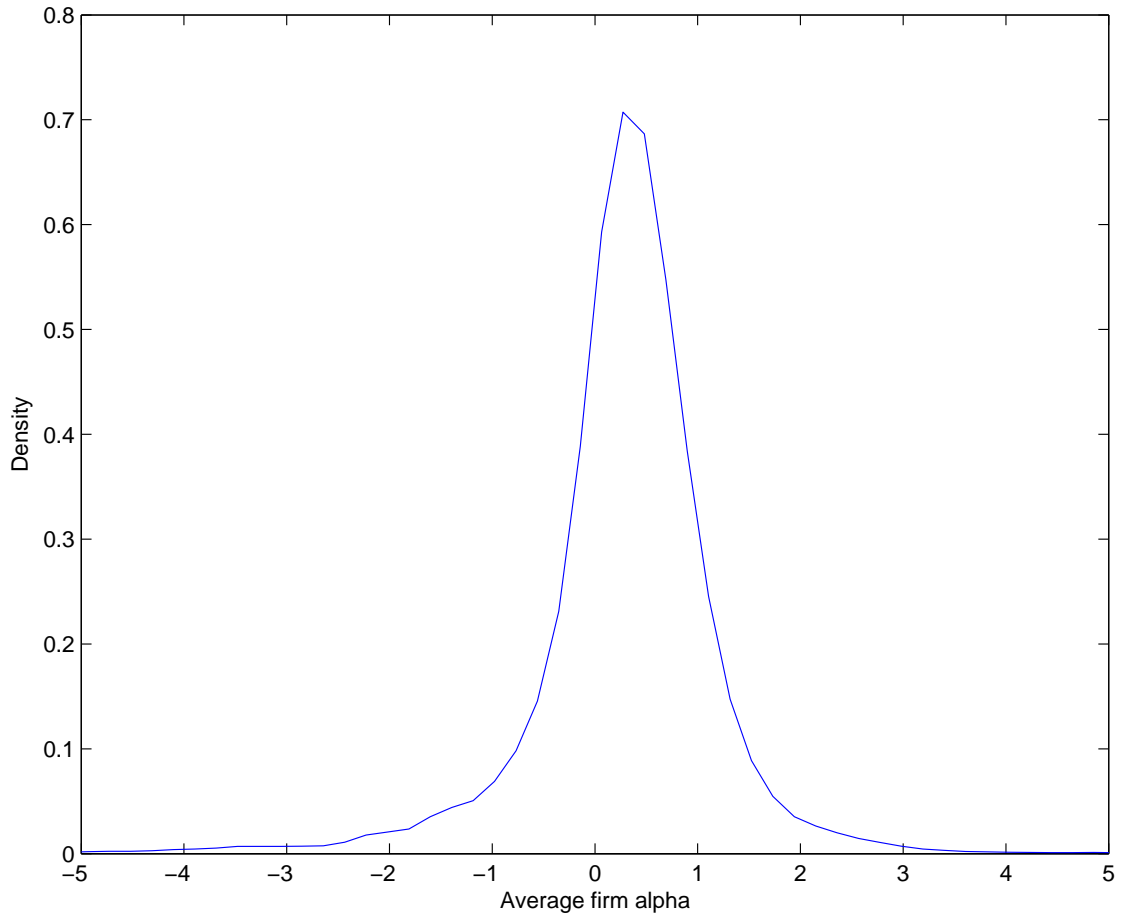
Figure 5: **Cross-Sectional Variation in Firm Betas**

This figure shows the evolution through time of the cross-sectional standard deviation of conditional firm betas. We first calculate at each iteration of the Gibbs sampler the conditional beta for firm i at time t based on the model in equation (7). Subsequently, we compute for every firm and every date the average beta over all iterations (posterior mean). Shaded areas indicate NBER recession periods.
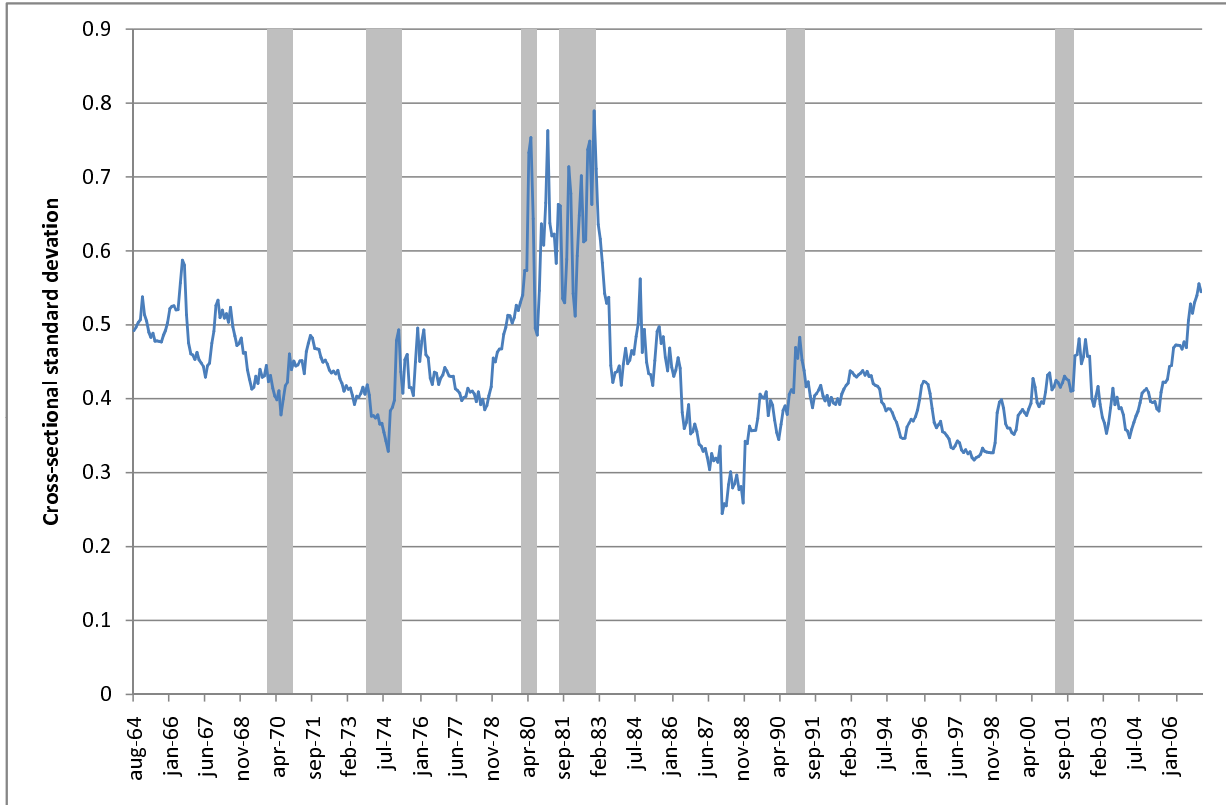
Figure 6: **Comparison of Cross-Sectional Variation in Firm Betas**

This figure compares the evolution through time of the cross-sectional standard deviation of different specifications for beta. Specifically, we compare the conditional beta given by equation (2), the realized beta in equation (3), the fundamentals-based beta in equation (5), and the static beta in the unconditional CAPM. Shaded areas indicate NBER recession periods.
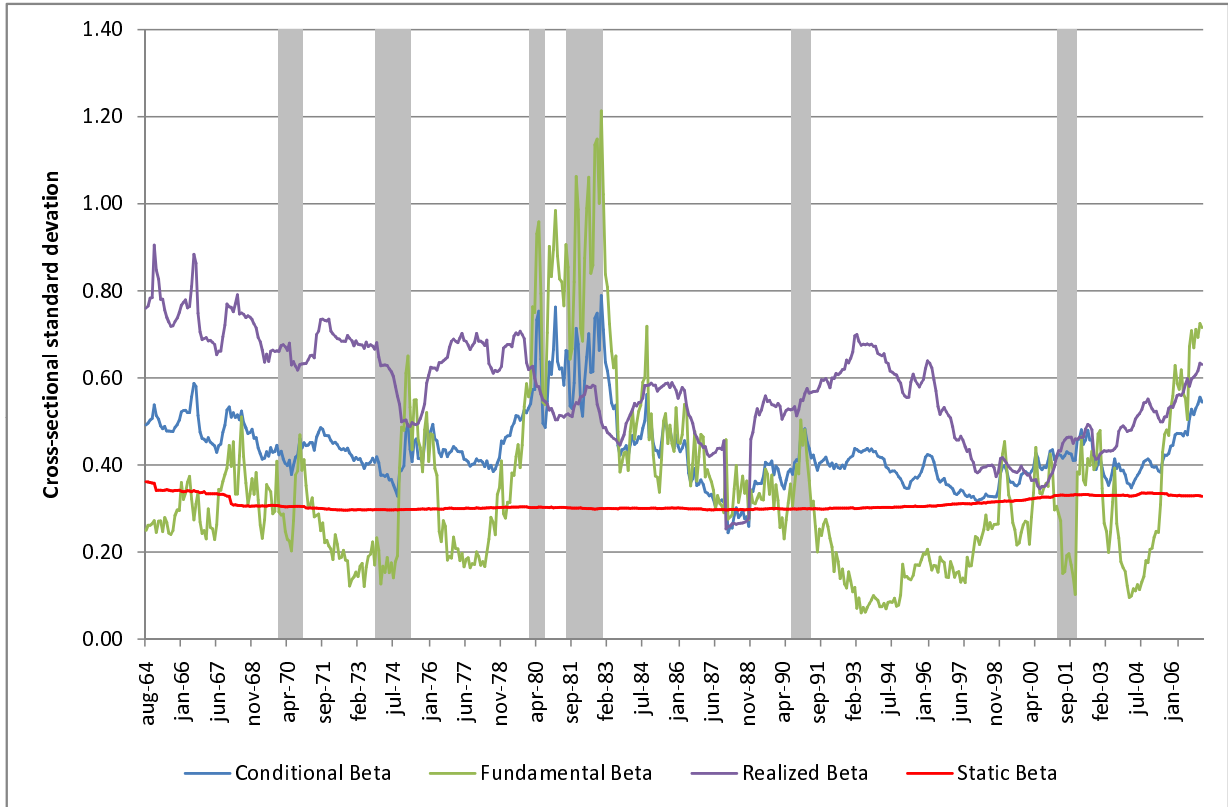
Figure 7: **Cross-Sectional Variation in Firm Alphas**

This figure shows the evolution through time of the cross-sectional standard deviation of conditional firm alphas. We first calculate at each iteration of the Gibbs sampler the conditional alpha for firm i at time t based on the model in equation (7). Subsequently, we compute for every firm and every date the average alpha over all iterations (posterior mean). Shaded areas indicate NBER recession periods.
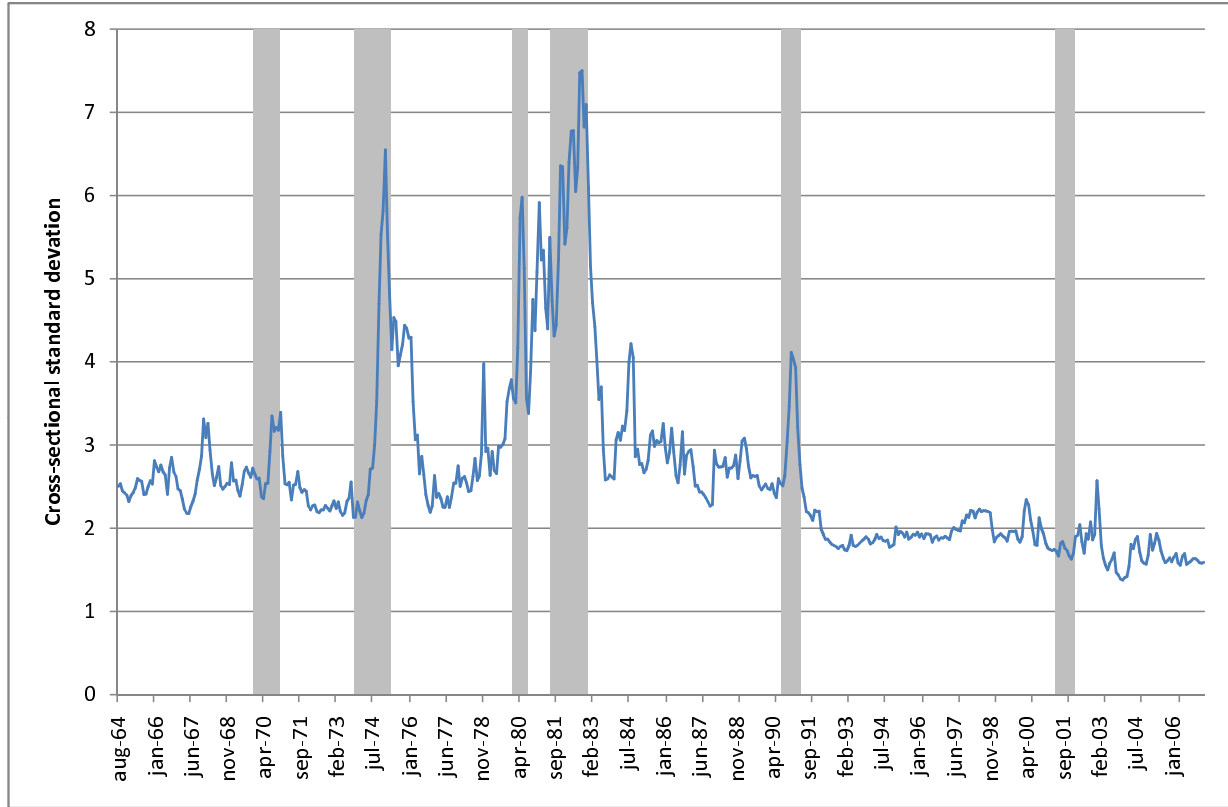
Figure 8: **Correlation between Firm Alphas and Betas**

This figure shows the evolution through time of the correlation between conditional firm alphas and betas. We first calculate at each iteration of the Gibbs sampler the conditional alpha and beta for firm i at time t based on the model in equation (7). Subsequently, we compute for every firm and every date the average alpha and beta over all iterations (posterior mean). We then calculate for every period the cross-sectional correlation between these alphas and betas. Shaded areas indicate NBER recession periods.
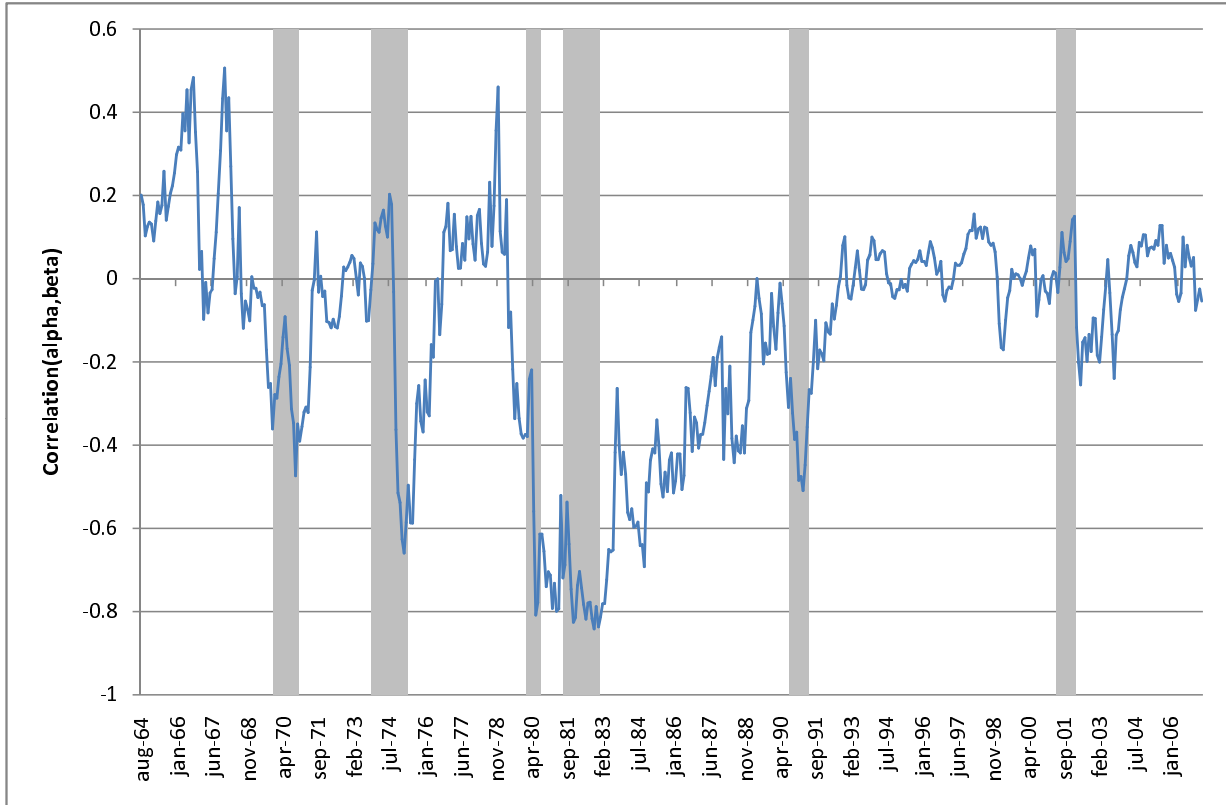
Figure 9: **Conditional Market Betas of Value and Growth Portfolios**

This figure plots the evolution over time of the conditional market beta for value and growth portfolios. We first construct 25 size-B/M portfolios by sorting stocks independently into size and B/M quintiles at the end of June. The 25 portfolios are then formed as the intersections of these five size and B/M quintiles. Subsequently, we calculate for every portfolio j at every time t the equally weighted cross-sectional average of the betas of the stocks in the portfolio. Beta for the value portfolio is defined as the average conditional beta of the five portfolios with highest B/M (Sx/B5) and beta for the growth portfolio is the average conditional beta of the five portfolios with lowest B/M (Sx/B1).