Assessing Market Beta Estimates*

Petri Jylhä
Aalto University

Yuekun Liu
Alliance Manchester Business School

Matthijs Lof
Aalto University

December 15, 2025

Abstract

We introduce a new measure, γ, for evaluating the accuracy of market beta estimates without relying on realized betas or strong model assumptions. γ is derived from the covariance between estimated betas and returns and is proportional to the correlation between estimated and true betas. Monte Carlo simulations show that γ reliably ranks beta estimates across a wide range of asset pricing models and has high power in detecting even modest differences in accuracy. Empirically, we apply our method to U.S. stocks, portfolios, and actively managed equity mutual funds, revealing meaningful differences in estimation accuracy across commonly used beta estimators.

Keywords: Market beta, CAPM, estimation methods
JEL classification: C52, C58, G12

*We thank Mathijs Cosemans (discussant), Keyan Liu, Yifan Li, Shifan Yu, and conference and seminar participants at the AEA 2025 (San Francisco), NVFEB 2024 (Rotterdam), SEM 2024 (Atlanta), QFFE 2025 (Marseille), EC^2 (Lugano), Aalto University, Bank of Finland, Helsinki GSE, and the Oxford-Man institute for useful comments. All errors are our own. Petri Jylhä (petri.jylha@aalto.fi); Yuekun Liu (yuekun.liu@manchester.ac.uk); Matthijs Lof (matthijs.lof@aalto.fi)

1 Introduction

Market beta, the sensitivity of an asset's return to market movements, is fundamental to asset pricing, corporate valuation, portfolio management, and investment analysis. Since the seminal work of Markowitz (1952) and Sharpe (1964), market beta has remained central to both academic research and financial practice (Graham and Harvey, 2001, Décaire et al., 2024), notably in estimating firms' cost of equity (e.g.,Fama and French, 1997; Da et al., 2012; Eaton et al., 2025) and in testing risk-return tradeoffs under the CAPM and its extensions (e.g., Lewellen and Nagel, 2006; Guo et al., 2017).

Despite its importance, how market beta should be estimated remains an open question. Empirical studies typically use historical return data and employ methods such as rolling-window OLS regressions (Fama and MacBeth, 1973), shrinkage toward the cross-sectional mean (Vasicek, 1973), adjustments for extreme returns (Welch, 2022), or corrections for asynchronous trading (Scholes and Williams, 1977, and Dimson, 1979). Different methods produce differing beta estimates, and evaluating their accuracy is challenging as true betas are unobservable.

One common evaluation method benchmarks estimated betas against ex-post realized betas (e.g., Hollstein and Prokopczuk, 2016, Cosemans et al., 2016, and Welch, 2022), implicitly assuming that realized betas closely approximate true betas. This assumption, however, is untestable, as realized betas themselves are estimated and subject to measurement error (cf. Patton, 2011, for a similar concern in volatility forecasting). An alternative method tests whether estimated market betas correlate with expected returns in the cross-section (e.g., Cosemans et al., 2016). Yet, a positive correlation with expected returns is neither a necessary nor sufficient condition for an accurate estimate of true beta: substantial evidence rejects the CAPM model and the pricing of market risk. Moreover, an estimated beta may correlate with expected returns because it is correlated with another priced factor even if it poorly reflects true beta.

In this paper, we introduce a novel and model-agnostic framework for evaluating market beta estimates directly against the unobserved true betas. Our framework flexibly nests a wide range of asset pricing models and does not require market beta to be a determinant of expected returns. Specifically, we propose a novel measure, γ, which is proportional to the cross-sectional correlation between estimated market betas and unobserved true betas. While the proportionality factor, the cross-sectional standard deviation of true market beta, is unobservable, it cancels out in pairwise comparisons. Therefore, γ provides a valid basis for ranking market beta estimators by accuracy. Estimation of γ is based on a simple time-series regression in which the cross-sectional covariance of estimated betas and returns, scaled by the cross-sectional standard deviation of estimated betas, is regressed on the market return. It is straightforward to test for statistically significant differences in γ between different beta estimators.

To validate our proposed measure, we conduct a battery of Monte Carlo simulations. We first generate stock return data with known time-varying true betas and construct noisy beta estimates. The data generating process consists of various settings with priced or unpriced market risk and potential correlation with other priced or unpriced factors. We show that our measure, γ, closely tracks the actual correlation between the constructed beta estimates and the true betas. Next, we compare pairs of noisy beta estimates and demonstrate that our test procedure exhibits high power in detecting even modest differences in estimation accuracy.

These simulations confirm that our framework reliably ranks different beta estimators under a variety of asset pricing conditions with strong discriminatory power.

We start our empirical analyses by applying our framework to data on U.S. common stocks from the CRSP database between 1978 and 2023. Our results offer several compelling insights. First, betas estimated from daily data consistently produce higher γ values than those based on monthly returns. This suggests that the use of daily data improves estimation accuracy. Second, widely used adjustments to OLS estimates, including shrinkage toward the cross-sectional mean (Vasicek, 1973), winsorization of outlier returns (Welch, 2022), exponential weighting (Boons, 2016), and corrections for non-synchronous trading (Scholes and Williams, 1977, and Dimson, 1979), consistently enhance beta estimation accuracy compared to the standard OLS benchmark. Notably, we find the estimator proposed by Frazzini and Pedersen (2014) to exhibit the highest accuracy among regression-based methods. We also find strong benefits from model averaging: the simple average across all regression-based beta estimates delivers a γ value that is significantly higher than any individual estimator. This suggests that averaging high-quality estimators further reduces idiosyncratic estimation noise and improves estimation accuracy.

Our analysis further extends beyond regression-based methods by incorporating option-implied betas (Buss and Vilkov, 2012). We find that betas implied by options with 182 days to maturity achieves the highest γ, indicating the most accurate estimate. However, differences in γ across option-implied betas of varying maturities are minimal beyond the 91-day maturity, consistent with their high correlations.

A common practice in cross-sectional asset pricing is to estimate betas from portfolios and assign them to individual stocks (e.g., Fama and French, 1992). We thus turn to examine whether portfolio-based betas improve accuracy relative to stock-level estimates. Our results show that portfolio-assigned betas consistently yield lower γ values than stock-level OLS betas, especially when the number of portfolios is small. Although increasing the number of portfolios narrows the gap, we find no setting in which portfolio-based betas significantly outperform individual stock betas.

Importantly, we show that the rankings of beta accuracy based on γ diverge markedly from those based on traditional evaluation metrics benchmarking estimated betas against realized beta, such as root mean squared error (RMSE) or correlation. For example, among adjusted regression-based estimators, the correlation between RMSE-based and γ-based rankings is only -15.5%, while RMSE and correlation rankings are nearly identical. These results suggest that differences across beta estimators are driven by estimation noise, which is precisely the dimension that our γ measure is designed to capture.

Finally, we extend our analysis to actively managed U.S. mutual funds. The mutual fund literature commonly uses lagged betas as proxies for funds' expected risk exposures in constructing predictive alpha (e.g., Carhart, 1997; Amihud and Goyenko, 2013; Cederburg et al., 2018). Our evidence suggests that explicitly correcting for non-synchronous trading using Dimson's (1979) approach improves beta accuracy relative to standard estimators.

Overall, our paper contributes to the literature by proposing a robust and broadly applicable measure for evaluating market beta accuracy. With our direct and model-agnostic framework, we enable researchers to make more informed choices in estimating market beta. The remainder of the paper is organized as follows: Section 2 introduces the γ framework; Section 3 validates the approach through Monte Carlo simulations; Section 4 presents empirical results using U.S. stocks and actively managed equity mutual funds; Section 5 concludes.

2 Methodology

2.1 Framework

We start from a simple decomposition of excess returns:

$$R_{i,t} = \beta_{i,t} R_{m,t} + e_{i,t}, \tag{1}$$

where $R_{i,t}$ is the excess return of asset i in time period t, $R_{m,t}$ is the excess return of the value-weighted market portfolio, and $e_{i,t}$ captures all other systematic and idiosyncratic components of $R_{i,t}$. The *unobservable* characteristic $\beta_{i,t}$ captures the asset's exposure to the market portfolio and is defined as:

$$\beta_{i,t} = \frac{E\left[R_{i,t} R_{m,t}\right] - E\left[R_{i,t}\right] E\left[R_{m,t}\right]}{E\left[R_{m,t}^2\right] - E\left[R_{m,t}\right]^2}. \tag{2}$$

Note that $\beta_{i,t}$ is time-varying and refers to the single-factor market beta, measuring the sensitivity of the asset's returns to market returns. While we focus on the single-factor market beta, we do not require market risk to be priced, nor do we rule out the possibility that other factors may also be priced. Specifically, the component $e_{i,t}$ in (1) can nest all other sources of variation in excess returns, including alpha, idiosyncratic returns, exposure to other factors, and microstructure noise. We are thus not imposing a specific asset pricing model, such as the CAPM. For any given $\{R_{i,t}, \beta_{i,t}, R_{m,t}\}$, there exists an $e_{i,t}$ such that (1) holds, implying that the equation provides a general decomposition compatible with any asset pricing model. For our simulations in Section 3, we consider various alternative specifications of $e_{i,t}$.

Let $\widetilde{\beta}_{i,t}$ denote an estimated beta for asset i in period t, where the panel of estimates is not required to be balanced. For each time period $t \in 1, \ldots, T$, let n_t denote the number of available cross-sectional observations. We remain agnostic about the beta estimation method or data used to estimate $\widetilde{\beta}_{i,t}$ and focus instead on evaluating the accuracy of these estimates relative to the true, unobserved beta $\beta_{i,t}$. Using the decomposition (1), the cross-sectional covariance between the estimated betas and excess returns can be expressed as:

$$Cov_t^{cs}\left[\widetilde{\beta}_{i,t}, R_{i,t}\right] = Cov_t^{cs}\left[\widetilde{\beta}_{i,t}, \beta_{i,t}\right] R_{m,t} + Cov_t^{cs}\left[\widetilde{\beta}_{i,t}, e_{i,t}\right], \tag{3}$$

in which $Cov_t^{cs}[.]$ refers to the cross-sectional covariance in period t. Dividing both sides by the cross-sectional standard deviation of estimated betas gives:

$$\frac{Cov_t^{cs}\left[\widetilde{\beta}_{i,t}, R_{i,t}\right]}{\sigma_t^{cs}\left[\widetilde{\beta}_{i,t}\right]} = \frac{Cov_t^{cs}\left[\widetilde{\beta}_{i,t}, \beta_{i,t}\right]}{\sigma_t^{cs}\left[\widetilde{\beta}_{i,t}\right]} R_{m,t} + \frac{Cov_t^{cs}\left[\widetilde{\beta}_{i,t}, e_{i,t}\right]}{\sigma_t^{cs}\left[\widetilde{\beta}_{i,t}\right]}. \tag{4}$$

Defining γ_t as the coefficient on $R_{m,t}$ in (4), that is, $\gamma_t \equiv \frac{Cov_t^{cs}\left[\widetilde{\beta}_{i,t}, \beta_{i,t}\right]}{\sigma_t^{cs}\left[\widetilde{\beta}_{i,t}\right]}$, it is straightforward to see that γ_t is proportional to the cross-sectional correlation between the estimated and true betas:

$$\gamma_t \propto \rho_t^{cs}\left[\widetilde{\beta}_{i,t}, \beta_{i,t}\right] \equiv \frac{Cov_t^{cs}\left[\widetilde{\beta}_{i,t}, \beta_{i,t}\right]}{\sigma_t^{cs}\left[\widetilde{\beta}_{i,t}\right] \sigma_t^{cs}\left[\beta_{i,t}\right]}, \tag{5}$$

or equivalently, $\gamma_t = \rho_t^{cs}\left[\widetilde{\beta}_{i,t}, \beta_{i,t}\right] \times \sigma_t^{cs}\left[\beta_{i,t}\right]$. When comparing two sets of beta estimates $\widetilde{\beta}_{i,t}^{(1)}$ and $\widetilde{\beta}_{i,t}^{(2)}$, applied to the same cross-section of assets, the cross-sectional standard deviation of the true beta $\sigma_t^{cs}[\beta_{i,t}]$ is identical. This implies that the set of beta estimates with the higher γ_t has a higher cross-sectional correlation with true beta.

While γ_t is not observable, its time-series mean γ can be inferred from a simple time-series regression. Given observable time-series of cross-sectional sample covariances $Cov_t^{cs}\left[\widetilde{\beta}_{i,t}, R_{i,t}\right]$, cross-sectional sample standard deviations $\sigma_t^{cs}\left[\widetilde{\beta}_{i,t}\right]$, and the market returns $R_{m,t}$, we estimate γ from the following time-series regression:

$$\frac{Cov_t^{cs}\left[\widetilde{\beta}_{i,t}, R_{i,t}\right]}{\sigma_t^{cs}\left[\widetilde{\beta}_{i,t}\right]} = \delta + \gamma R_{m,t} + \xi_t. \tag{6}$$

The OLS estimate of γ is our main measure of interest. It represents the time-series mean of γ_t, which is proportional to the cross-sectional correlation between the estimated and true betas.

Estimating γ by OLS is unbiased and consistent if ξ_t and $R_{m,t}$ are uncorrelated. As can be seen from (4), this condition is satisfied when the market return $R_{m,t}$ is uncorrelated with both $Cov_t^{cs}\left[\widetilde{\beta}_{i,t}, e_{i,t}\right]$ and $\sigma_t^{cs}\left[\widetilde{\beta}_{i,t}\right]$. In Appendix A, we derive that the correlation between the cross-sectional covariance $Cov_t^{cs}\left[\widetilde{\beta}_{i,t}, e_{i,t}\right]$ and $R_{m,t}$ is zero if the estimated stock-level betas $\widetilde{\beta}_{i,t}$ are uncorrelated with $R_{m,t}$. It is important to note that $\widetilde{\beta}_{i,t}$ is estimated from historical data, such that the estimated beta of stock i in period t is based on data prior to period t. Given that market returns are notoriously hard to predict using lagged information (e.g., Welch and Goyal, 2008), it is unlikely that $\widetilde{\beta}_{i,t}$ or its cross-sectional standard deviation $\sigma_t^{cs}\left[\widetilde{\beta}_{i,t}\right]$ have any significant correlation with $R_{m,t}$. In Appendix A, we provide evidence consistent with the absence of such correlation, for a wide range of beta estimators, to alleviate any concerns about estimation bias in γ.

We can also test directly for differences in estimation accuracy between two sets of beta estimates $\widetilde{\beta}_{i,t}^{(1)}$ and $\widetilde{\beta}_{i,t}^{(2)}$, using the following regression:

$$\frac{Cov_t^{cs}\left[\widetilde{\beta}_{i,t}^{(2)}, R_{i,t}\right]}{\sigma_t^{cs}\left[\widetilde{\beta}_{i,t}^{(2)}\right]} - \frac{Cov_t^{cs}\left[\widetilde{\beta}_{i,t}^{(1)}, R_{i,t}\right]}{\sigma_t^{cs}\left[\widetilde{\beta}_{i,t}^{(1)}\right]} = \Delta\delta + \Delta\gamma R_{m,t} + \Delta\xi_t. \tag{7}$$

The slope coefficient $\Delta\gamma$ in (7) captures the difference between γ estimates for $\widetilde{\beta}_{i,t}^{(2)}$ and $\widetilde{\beta}_{i,t}^{(1)}$, as obtained from (6). A positive value of $\Delta\gamma$ implies that $\widetilde{\beta}_{i,t}^{(2)}$ has a higher correlation with true beta than $\widetilde{\beta}_{i,t}^{(1)}$: $\rho^{cs}\left[\widetilde{\beta}_{i,t}^{(2)}, \beta_{i,t}\right] > \rho^{cs}\left[\widetilde{\beta}_{i,t}^{(1)}, \beta_{i,t}\right]$.

The advantage of estimating $\Delta\gamma$ in a single regression is that it enables straightforward statistical inference using a standard t-test. This allows us to formally test whether one set of beta estimates is more closely correlated with true beta than another, on average over the sample period. In the next section, we examine how the structure of estimation error affects the interpretation of γ.

2.2 Noise and bias in beta estimates

An estimator of beta can be inaccurate either because of estimation noise or estimation bias. Consider an unbiased estimator, where the only deviation from the true beta is noise:

$$\widetilde{\beta}_{i,t} = \beta_{i,t} + \nu_{i,t}, \tag{8}$$

where $\nu_{i,t}$ is *i.i.d.* zero-mean noise. The correlation between $\beta_{i,t}$ and $\widetilde{\beta}_{i,t}$ depends only on the relative magnitudes of $Var[\beta]$ and $Var[\nu]$:

$$\rho\left[\widetilde{\beta}_{i,t}, \beta_{i,t}\right] = \sqrt{\frac{Var[\beta]}{Var[\beta] + Var[\nu]}}. \tag{9}$$

Consider two panels of unbiased estimates $\widetilde{\beta}_{i,t}^{(1)}$ and $\widetilde{\beta}_{i,t}^{(2)}$, which are both defined as in (8), but differ in the variance of noise $Var[\nu]$. Sorting $\widetilde{\beta}_{i,t}^{(1)}$ and $\widetilde{\beta}_{i,t}^{(2)}$ on their correlation with true beta is thus equivalent to (reverse) sorting on the variance of $\nu_{i,t}$: the estimator with a lower variance of noise $Var[\nu]$ will have a higher correlation with true beta and therefore a higher value of γ.

In addition to noise, an estimator $\widetilde{\beta}_{i,t}$ can also be biased. Assume that the estimator can be represented as the following linear transformation of true beta:

$$\widetilde{\beta}_{i,t} = (1+\delta)\left(\beta_{i,t} + \nu_{i,t}\right) - \delta, \tag{10}$$

where δ is a bias term. When $\delta = 0$, $\widetilde{\beta}_{i,t}$ reduces to the unbiased case in (8). When $\delta > 0$, $\widetilde{\beta}_{i,t}$ overestimates high-beta stocks ($\beta_{i,t} > 1$) and underestimates low-beta stocks ($\beta_{i,t} < 1$), with the direction of bias reversing when $\delta < 0$. Crucially, even when $\delta \neq 0$, the correlation between $\widetilde{\beta}_{i,t}$ and $\beta_{i,t}$ remains unchanged from (9). As a result, our γ measure reflects only the impact of the noise component in (10) and is unaffected by bias. Two sets of beta estimates with identical noise variance but different values of δ will therefore have the same correlation with true beta, and hence the same value of γ. This implies that γ is a measure of cross-sectional accuracy when the objective is to *sort* assets by their true beta and provides effective separation of high-beta from low-beta assets, even if the estimated betas are biased in level. In section 4.6, we show empirically that, among commonly used beta estimators, noise rather than bias is the dominant source of estimation error.

A potentially appealing approach to evaluate the bias in a panel of beta estimates $\widetilde{\beta}_{i,t}$ is to consider the hypothetical regression of true betas on estimated betas:

$$\beta_{i,t} = a + b\widetilde{\beta}_{i,t} + \varepsilon^{\beta}_{i,t}. \tag{11}$$

The OLS estimate of the slope coefficient, which is equal to $\widehat{b} = \frac{Cov\left[\widetilde{\beta}_{i,t}, \beta_{i,t}\right]}{Var\left[\widetilde{\beta}_{i,t}\right]}$, cannot be obtained directly without observing true beta. However, b could be estimated in a similar way as we estimate γ, by replacing the denominator $\sigma^{cs}_t\left[\widetilde{\beta}_{i,t}\right]$ in (4) by the variance of estimated betas. This gives:

$$\frac{Cov^{cs}_t\left[\widetilde{\beta}_{i,t}, R_{i,t}\right]}{Var^{cs}_t\left[\widetilde{\beta}_{i,t}\right]} = \frac{Cov^{cs}_t\left[\widetilde{\beta}_{i,t}, \beta_{i,t}\right]}{Var^{cs}_t\left[\widetilde{\beta}_{i,t}\right]} R_{m,t} + \frac{Cov^{cs}_t\left[\widetilde{\beta}_{i,t}, e_{i,t}\right]}{Var^{cs}_t\left[\widetilde{\beta}_{i,t}\right]}, \tag{12}$$

such that a regression of $\frac{Cov^{cs}_t\left[\widetilde{\beta}_{i,t}, R_{i,t}\right]}{Var^{cs}_t\left[\widetilde{\beta}_{i,t}\right]}$ on $R_{m,t}$ provides a slope coefficient that can be interpreted as b, similar to our regression (6) to estimate γ.

Hollstein and Prokopczuk (2016) apply the regression model (11) directly, by substituting the true beta with a measure of realized beta, and consider the hypothesis $H_0 : b = 1$ as a test of unbiasedness of the estimated betas. Based on our implied estimation method, we can test this hypothesis without requiring an estimate of realized beta. However, the hypothesis $H_0 : b = 1$ is ambiguous as to whether $\widetilde{\beta}_{i,t}$ is a good estimate of $\beta_{i,t}$. To see this, assume that the estimated beta is linearly related to true beta as in (10), leading to the OLS estimator

$\widehat{b}$ representing:

$$\widehat{b} = \frac{Cov\left[\widetilde{\beta}_{i,t}, \beta_{i,t}\right]}{Var\left[\widetilde{\beta}_{i,t}\right]} = \frac{(1+\delta)\, Var\left[\beta_{i,t}\right]}{(1+\delta)^2 \left(Var\left[\beta_{i,t}\right] + Var\left[\nu_t\right]\right)}. \tag{13}$$

It is straightforward to see that if $\delta = 0$ (i.e. $\widetilde{\beta}_{i,t}$ is unbiased), then $\widehat{b} \leq 1$, and an estimate of $\widehat{b}$ close to one implies that $\widetilde{\beta}_{i,t}$ is a precise estimator with small variance of estimation noise $\nu_{i,t}$. However, when $\delta \neq 1$, it is possible for $\widehat{b}$ to be above, below, or equal to one. The restriction $b = 1$ is therefore neither a necessary nor sufficient condition for $\widetilde{\beta}_{i,t}$ to be unbiased and is ambiguous about the precision of $\widetilde{\beta}_{i,t}$.[1] For this reason, we do not focus on b, but rather on γ as our accuracy measure of interest. A higher γ implies a higher correlation between $\widetilde{\beta}_{i,t}$ and $\beta_{i,t}$, without requiring any assumptions on the unbiasedness or other properties of $\widetilde{\beta}_{i,t}$.

A more complete measure of estimation accuracy is the Mean Squared Error (MSE):

$$MSE_t^{cs}\left[\widetilde{\beta}_{i,t}\right] = n_t^{-1} \sum_{i=1}^{n_t} \left(\widetilde{\beta}_{i,t} - \beta_{i,t}\right)^2. \tag{14}$$

Assuming again that the estimated beta is a linear function of true beta as in (10), we get $MSE_t^{cs}\left[\widetilde{\beta}_{i,t}\right] = n_t^{-1} \sum_{i=1}^{n_t} \left(\delta\left(\beta_{i,t} - 1\right) + (1+\delta)\, \nu_{i,t}\right)^2$, which shows that the MSE captures both the magnitude of bias (through δ) and noise (through ν) in the beta estimate $\widetilde{\beta}$. Hollstein and Prokopczuk (2016) and Hollstein et al. (2020) use the MSE to evaluate the precision of beta estimates against a realized beta benchmark. Since true betas are unobservable, computing the MSE requires a proxy such as realized beta. As we show in Section 4.6, rankings of beta estimates based on their MSE with respect to realized beta closely

[1] Hollstein and Prokopczuk (2016) consider in fact the *joint* hypothesis $H_0 : a = 0; b = 1$. Our point about the ambiguity of the restriction $b = 1$ does not depend on the level of a.

align with those based on their correlation with realized beta. In contrast, these rankings diverge substantially from those based on the γ measure. Our evidence suggests that bias is a relatively minor component of estimation error, and that differences are primarily driven by noise.

3 Simulations

We validate our method using Monte Carlo simulations, where the true betas are known by construction. This allows us to directly assess the cross-sectional correlation between simulated beta estimates and the true betas. We generate stock-level returns using a known data-generating process, simulate pseudo-estimated betas by applying controlled levels of noise and bias, and compute the corresponding γ values. This simulation framework allows us not only to validate whether the estimated γ reliably tracks the cross-sectional correlation between estimated and true betas, but also to assess the statistical power of our procedure in detecting differences in estimation accuracy across beta measures.

We start by simulating true betas for 500 assets such that each asset's beta is time-varying but persistent while making sure that the cross-sectional average of betas in each period is one and their standard deviation is 0.5.[2] Next, we simulate 240 months of returns for each asset following the return-generating process in equation (1), varying the assumptions about the error term $e_{i,t}$. We then simulate pseudo-estimated betas, $\widetilde{\beta}_i$, from equation (10), where the bias term δ is drawn from a uniform distribution: $\delta \sim \mathcal{U}(-0.25,\ 0.25)$, and the standard

[2]Technically, the betas follow an AR(1) process with the autocorrelation parameter set to 0.95. The results presented here are not sensitive to the level of autocorrelation of betas.

deviation of the noise term ν_i is drawn from $\mathcal{U}(0.0,\ 0.5)$.[3] For each simulation, we estimate the γ coefficient using equation (6), and also compute the direct correlation between β_i and $\widetilde{\beta}_i$, denoted ρ. This yields one $(\gamma,\ \rho)$ pair. Repeating the procedure 100 times generates 100 such pairs plotted in Figure 1.

The simulation runs differ in how accurately the pseudo-estimated betas match the true betas. In some runs, the bias term δ is close to zero and the noise term ν has low variance, resulting in a high correlation between $\widetilde{\beta}_i$ and β_i. In these cases, we expect a high value of the estimated γ coefficient. In contrast, when δ deviates significantly from zero and the variance of ν is high, the pseudo-estimates are less accurate, which results in a lower value of ρ and correspondingly a lower γ. If our method reliably captures cross-sectional accuracy, the simulated $(\gamma,\ \rho)$ pairs in Figure 1 should lie along a positively sloped line, consistent with equation (5).

<FIGURE 1 HERE>

In Panel A of Figure 1, we assume that the CAPM is the true asset pricing model. We implement this by assuming that the error term $e_{i,t}$ in (1) only contains an idiosyncratic term, $e_{i,t} = \epsilon_{i,t}$, where $\epsilon_{i,t}$ is i.i.d. Normal. In Panel B, we assume a flat security market line, where expected returns do not depend on beta. This is equivalent to defining $e_{i,t} = (1 - \beta_{i,t})E\left[R_{m,t}\right] + \epsilon_{i,t}$. In panels C through F, we introduce an additional factor, F_t, into the return generating process, such that $e_{i,t} = \theta_{i,t}F_t + \epsilon_{i,t}$. In all four cases, market risk remains priced, but we vary the mean of F_t and the correlation between the factor exposures

[3] Note that we do not estimate the betas using the simulated return data and some arbitrarily chosen estimation method. Rather, we generate the estimated betas as a function of the true beta, a noise term, and a bias term.

β_i and θ_i. In Panel C, F_t has a zero mean and $Corr(\beta_i,\ \theta_i) = 0$. In Panel D, F_t has a non-zero mean, equal to half of the mean of $R_{m,t}$, and $Corr(\beta_i,\ \theta_i) = 0$. In Panel E, F_t has a zero mean and $Corr(\beta_i,\ \theta_i) = 0.2$. Finally, in Panel F, F_t has the same non-zero mean and $Corr(\beta_i,\ \theta_i) = 0.2$.

Across all panels, the varying assumptions regarding the return-generating process produce a remarkably consistent, linearly increasing relation between the estimated γ and the direct ρ. This consistency suggests that our method reliably ranks beta estimates based on their correlation with true beta, regardless of the underlying asset pricing model. Recall that $\gamma = \rho[\beta, \widetilde{\beta}] \times \sigma[\beta]$, which implies that when the pseudo-estimated betas are highly correlated with the true betas, the estimated γ should approach 0.5, the standard deviation of β_i in the simulation design.

In Figure 2, we assess the statistical power of the testing procedure described in (7). Like above, we draw a panel of 240 months of returns for 500 assets with randomly generated time-varying true betas. We then construct two sets of pseudo-estimated betas, $\widetilde{\beta}_i^{(1)}$ and $\widetilde{\beta}_i^{(2)}$, following equation (10). These beta estimates differ in terms of their correlation with the true betas. We assume both estimates are unbiased by setting $\delta = 0$. We assign a standard deviation of 0.1 to the noise term ν_i for $\widetilde{\beta}_i^{(1)}$, resulting in a correlation of 0.98 with the true beta. For $\widetilde{\beta}_i^{(2)}$, we vary the standard deviation of ν_i between 0.1 and 0.14, producing correlations with the true beta ranging from 0.96 to 0.98. We then test the hypothesis $H_0 : \Delta\gamma = 0$ as described in (7), evaluating whether the two sets of beta estimates differ significantly in their accuracy. For each level of correlation between β and $\widetilde{\beta}_i^{(2)}$, we repeat the procedure 10,000 times. Figure 2 displays the resulting power curves, rejection rates of the null at the 5% (blue) and the 1% (red) significance levels, as a function of the difference in correlation. The six panels of Figure 2, like those in Figure 1, correspond to different assumptions about the return-generating process.

<FIGURE 2 HERE>

Three key findings arise from Figure 2. First, the test has appropriate size. When the two panels of estimated betas, $\widetilde{\beta}_i^{(1)}$ and $\widetilde{\beta}_i^{(2)}$, have the same correlation with the true beta, the test rejects the null hypothesis, $\Delta\gamma = 0$, 5% and 1% of the time at the 5% and 1% significance levels, respectively. Second, the power curves are steep, indicating the test can detect even modest differences in the accuracy of beta estimates. For example, in Panel A, when the correlation difference between $\widetilde{\beta}_i^{(1)}$ and $\widetilde{\beta}_i^{(2)}$ is only 0.005, the null hypothesis is rejected 50% of the time at the 5% level and 27% of the time at the 1% level. When the difference in correlation is 0.010 (0.015), the corresponding rejection rates rise to 96% and 87%, respectively (100% at both levels). This suggests that the testing procedure is very powerful. Third, the power curves are nearly identical across all six panels of Figure 2. This again implies that the performance of the test is robust to different assumptions about the true asset pricing model.

Overall, the Monte Carlo simulations provide strong validation for our approach. The γ measure reliably ranks competing beta estimates by their correlation with true beta. We also demonstrate that our testing framework delivers high power even when differences in accuracy are modest. Importantly, these results hold across a wide range of asset pricing models. This robustness highlights the empirical relevance of our method and motivates its application to evaluating the cross-sectional accuracy of beta estimates in the sections that follow.

4 Empirical results

In this section, we apply our method to real data by comparing multiple beta estimation approaches commonly used in the literature.

4.1 Stock data

We build our sample using all NYSE, AMEX, and NASDAQ listed common stocks with a CRSP share code of 10 or 11, from January 1973 to December 2023. For each stock, we obtain daily and monthly returns, prices, and shares outstanding from CRSP. We also obtain daily and monthly value-weighted market returns from CRSP. Daily and monthly risk-free rates are obtained from Kenneth French's data library. For a stock-month observation to be included in our sample, we apply the following filters: *(i)* at least 200 non-missing daily return observations in the preceding 12 calendar months, *(ii)* at least 750 non-missing daily return observations in the preceding 60 calendar months, and *(iii)* at least 48 non-missing monthly return observations in the preceding 60 months. This leads to a sample of 1,974,590 observations, spanning 552 months (January 1978 to December 2023) and 16,342 unique stocks.

4.2 OLS beta estimates

For each stock-month observation in our sample, we estimate betas $\widetilde{\beta}_{i,t}$ by ordinary least squares (OLS) using either daily or monthly return observations:[4]

[4]To estimate the market beta as defined in (2), we only consider beta estimates from a single-factor regression, rather than partial beta estimates from a multi-factor regression. As we discuss in Section 2.1 and demonstrate in our simulations, this does not imply that there are no other factors, nor that the single-factor model is the correct asset pricing model.

Daily OLS beta ($\widetilde{\beta}_{i,t}^{OLS_D^K}$): We estimate beta in month t for stock i by regressing the stock's daily excess return on the market excess return: $R_{i,\tau} = \alpha_{i,t} + \beta_{i,t} R_{m,\tau} + \varepsilon_{i,\tau}$. The rolling estimation window includes all daily observations over the K months prior to month t. We consider $K = 6$, 12, 24, 36, and 60 months, with $K = 12$ corresponding to the standard practice in the literature, as in Fama and MacBeth (1973).

Monthly OLS beta ($\widetilde{\beta}_{i,t}^{OLS_M^K}$): Estimated similar to $\widetilde{\beta}_{i,t}^{OLS_D^K}$, using rolling estimation windows of *monthly* stock and market returns over the K months prior to month t. We consider $K = 12$, $K = 24$, $K = 36$, and $K = 60$.

<TABLE 1 HERE>

Table 1 reports summary statistics of the betas estimated by OLS using different windows of daily and monthly returns. Panel A presents the distribution of the estimated betas. The first four rows report the time-series average of the equal-weighted and value-weighted cross-sectional average and cross-sectional standard deviation of $\widetilde{\beta}_{i,t}$, for each of the beta estimators. As expected, the value-weighted mean betas are all close to one. Cross-sectional standard deviations are around 0.6 for daily betas and notably larger for monthly betas, especially with shorter windows (e.g., 1.37 for OLS_M^{12}), reflecting estimation noise from small samples. Time-series standard deviations also decline with longer windows, indicating that beta estimates become more stable as the estimation horizon increases.

Panel B reports the average monthly cross-sectional correlations between different beta estimates. Daily betas are generally highly correlated with each other, while correlations fall as horizons diverge. A similar pattern holds for monthly betas, although their correlations are on average lower. Correlations between daily and monthly betas are notably lower, especially at shorter horizons. Overall, Table 1 shows that both return frequency and estimation horizon generate meaningful variation in beta estimates.

Given this sample of estimated betas and returns, we compute, for each panel of estimated betas, a time-series of its monthly sample covariance with returns, denoted as $Cov_t^{cs}\left[\widetilde{\beta}_{i,t}, R_{i,t}\right]$, and its corresponding cross-sectional standard deviation, $\widehat{\sigma}_t^{cs}\left[\widetilde{\beta}_{i,t}\right]$. Following Eq. (6), we estimate γ by regressing the ratio of the covariance to the standard deviation $\frac{Cov_t^{cs}\left[\widetilde{\beta}_{i,t}, R_{i,t}\right]}{\sigma_t^{cs}\left[\widetilde{\beta}_{i,t}\right]}$ on the value-weighted market return $R_{m,t}$. This procedure yields a single γ estimate for each panel of beta estimates.

<TABLE 2 HERE>

Table 2 presents γ estimates for a range of OLS betas from different frequencies and estimation windows. For the baseline estimate that uses a 12-month window using daily returns (OLS_D^{12}), the γ estimate is 0.34. According to equation (5), γ is proportional to the cross-sectional correlation between the estimated beta and true betas, scaled by the standard deviation of the true beta. While the scaling factor is unobservable, it is fixed across all estimated betas. Therefore, a higher value of γ directly implies more accurate beta estimates.

Among the daily OLS estimates, increasing the estimation window modestly improves performance. For example, the 24 and 36-month estimation window both achieve a positive $\Delta\gamma$, though the differences are not statistically significant. In contrast, shortening the window to 6-month reduces γ to 0.311, with a significant $\Delta\gamma = -0.028$. This suggests that extending the estimation window reduces noise, but with diminishing returns beyond a certain length. Beta estimates based on monthly returns are consistently less accurate, with γ ranging from 0.236 to 0.327. This likely reflects large estimation error due to a limited number of observations at lower frequency, as well as reduced responsiveness of long-window estimates, making them less effective at capturing time-varying risks.

Building on these results, we take the 24-month OLS beta estimated from daily returns (OLS_D^{24}) as our benchmark. In the following section, we compare it to a range of commonly used adjusted beta estimators, including shrinkage, winsorization, and alternative weighting approaches.

4.3 Adjusted regression-based beta estimates

The literature has proposed several adjustments to improve OLS-based beta estimates. One common approach is to apply shrinkage or winsorization to reduce the impact of estimation noise and extreme return observations. We first consider four such adjustments:

Blume beta ($\widetilde{\beta}_{i,t}^{BLM}$): To correct for historical data bias in the OLS beta, Blume (1971) proposes linear shrinkage of the OLS estimator. Following Bloomberg's methodology, we apply a shrinkage factor of $\frac{2}{3}$ to the benchmark OLS beta estimated from 24-month windows of daily data: $\widetilde{\beta}_{i,t}^{BLM} = \frac{1}{3} + \frac{2}{3}\widetilde{\beta}_{i,t}^{OLS}$.

Slope-Winsorized beta ($\widetilde{\beta}_{i,t}^{BSW}$ and $\widetilde{\beta}_{i,t}^{BSWA}$): Following Welch (2022), $\widetilde{\beta}_{i,t}^{BSW}$ is estimated by rolling-window regressions of daily winsorized stock returns on a constant and the daily market returns. Daily stock returns are winsorized by establishing minimum and maximum coefficient slopes relative to market returns. Specifically, the winsorized return for stock i on day τ, denoted $R_{sw,i,\tau}$, is adjusted to fall within the range: $R_{sw,i,\tau} \in [1.0-\Delta, 1.0+\Delta] \times R_{m,\tau}$, where Δ is set equal to three. Further building on Welch (2022), $\widetilde{\beta}_{i,t}^{BSWA}$ is estimated using the same winsorizing scheme and by weighted-least-squares to ensure decay of weights over time, with the latest observation having the largest weight. The weight of observation τ is set to $e^{-\frac{2}{256}\cdot(\bar{\tau}-\tau)}$, where $\bar{\tau}$ denotes the end of the estimation window. For comparison to the benchmark, we apply estimation windows of 24 months of daily returns.

Vasicek beta ($\widetilde{\beta}_{i,t}^{VCK}$): Following Vasicek (1973), we shrink the 24-month daily OLS beta towards the cross-sectional mean: $\widetilde{\beta}_{i,t}^{VCK} = w_{i,t}\widetilde{\beta}_{i,t}^{OLS} + (1-w_{i,t})\overline{\widetilde{\beta}_{t}^{OLS}}$, where $\overline{\widetilde{\beta}_{t}^{OLS}}$ is the cross-sectional mean of $\widetilde{\beta}^{OLS}$. The Bayesian shrinkage factor, $w_{i,t}$, is defined as $\frac{\sigma_t^2}{\sigma_{i,t}^2+\sigma_t^2}$, where σ_t^2 is the cross-sectional variance of OLS betas for all stocks in period t and $\sigma_{i,t}^2$ is the variance (squared standard error) of $\widetilde{\beta}_{OLS}$ for each stock.

Next, we consider Exponentially Weighted Least Squares (EWLS) estimation, which assigns greater weight to more recent observations within the estimation window.

Exponentially-weighted Least Squares Beta ($\widetilde{\beta}_{i,t}^{EWLS}$): Following Boons (2016) and Hollstein et al. (2020), we estimate beta by employing weighted least squares regression, using weights $\frac{\exp(-|\bar{\tau}-\tau|\times h)}{\sum_{\tau=1}^{\bar{\tau}}\exp(-|\bar{\tau}-\tau|\times h)}$. In this expression, $\bar{\tau}$ denotes the end of the estimation sample, and τ refers to the observation date. The decay parameter h is defined as $\frac{\log(2)}{l}$, with l set to either 84 days ($\widetilde{\beta}_{i,t}^{EWLS,84}$) or 168 days ($\widetilde{\beta}_{i,t}^{EWLS,168}$), following Hollstein et al. (2020). The rolling estimation window includes all daily observations over the 24 months prior to month t.

We then consider corrections that address nonsynchronous trading, particularly relevant for thinly traded stocks. We examine two such adjustments:

Dimson Beta ($\widetilde{\beta}_{i,t}^{DIM}$): Dimson (1979) introduces both current and lagged market returns in the regressions to mitigate the effects of nonsynchronous trading. We follow the specification by Lewellen and Nagel (2006), including four daily lagged market returns: $R_{i,\tau} = \alpha_{i,t} + \beta_{0,i,t} R_{m,\tau} + \beta_{1,i,t} R_{m,\tau-1} + \beta_{2,i,t} \frac{1}{3} \left(R_{m,\tau-2} + R_{m,\tau-3} + R_{m,\tau-4} \right) + \varepsilon_{i,\tau}$. The Dimson beta is defined as $\widetilde{\beta}_{i,t}^{DIM} = \sum_{j=0}^{2} \beta_{j,i,t}$. The rolling estimation window includes all daily observations over the 24 months prior to month t.

Scholes-Williams Beta ($\widetilde{\beta}_{i,t}^{SW}$): Following Scholes and Williams (1977), we conduct for each stock i in month t three separate regressions of the stock's excess returns on contemporaneous, lagged, and lead market excess returns. The three betas are defined as: contemporaneous beta, $\beta_{i,t} = \frac{Cov[R_{i\tau}, R_{m,\tau}]}{Var[R_{m,\tau}]}$; lagged beta, $\beta_{i,t}^{-} = \frac{Cov[R_{i\tau}, R_{m,\tau-1}]}{Var[R_{m,\tau-1}]}$; and leading beta, $\beta_{i,t}^{+} = \frac{Cov[R_{i\tau}, R_{m,\tau+1}]}{Var[R_{m,\tau+1}]}$. The composite Scholes-Williams beta is defined as $\widetilde{\beta}_{i,t}^{SW} = \frac{\beta_{i,t}^{-} + \beta_{i,t} + \beta_{i,t}^{+}}{1 + 2\rho_{m,t}}$, where $\rho_{m,t}$ is the autocorrelation coefficient of the market excess return, $\rho_{m,t} = \frac{Cov[R_{m,\tau}, R_{m,\tau-1}]}{\sigma[R_{m,\tau}]\sigma[R_{m,\tau-1}]}$. The rolling estimation window includes all daily observations over the 24 months prior to month t.

Finally, we consider two beta estimators that are not direct adjustments of the OLS beta:

Frazzini-Pedersen Beta ($\widetilde{\beta}_{i,t}^{FP}$): Following Frazzini and Pedersen (2014), we calculate beta as $\widetilde{\beta}_{i,t}^{FP} = \rho_{i,t}\frac{\sigma_{i,t}}{\sigma_{M,t}}$, with $\sigma_{i,t}$ and $\sigma_{M,t}$ denoting the estimated volatility of stock returns and market returns, respectively, and $\rho_{i,t}$ denoting the estimated correlation coefficient between market and stock returns. Volatilities are estimated using one-year rolling windows of daily log returns. Correlations are instead estimated over a five-year window, using overlapping three-day log returns: $r_{i,\tau}^{3d} = \sum_{k=0}^{2} \ln(1 + r_{i,\tau+k})$.

Average Beta ($\widetilde{\beta}_{i,t}^{AVG}$): The average regression-based beta of stock i in month t is the simple equal-weighted average of the 10 betas estimated from daily returns described above (including the baseline 24-month OLS).

<TABLE 3 HERE>

Table 3 provides summary statistics and correlations for the adjusted regression-based betas. For comparison, we also report a measure of **realized beta** ($\widetilde{\beta}^{RB}$). Realized beta for stock i in month t is estimated in-sample, using the daily returns observed during month t, in contrast to regression-based betas that rely on historical return windows prior to month t. Following Andersen et al. (2006), realized beta is estimated as:

$$\widetilde{\beta}_{i,t}^{RB} = \frac{\sum_{\tau=1}^{T_t} r_{i,\tau} r_{m,\tau}}{\sum_{\tau=1}^{T_t} r_{m,\tau}^2}, \tag{15}$$

where $r_{i,\tau} = log\,(1 + R_{i,\tau})$ is the log-return of stock i on day τ and $r_{m,\tau} = log\,(1 + R_{m,\tau})$ is the log-return of the market portfolio. T_t refers to the number of trading days in month t.

Panel A shows that most beta estimators generate similar average betas, especially under value-weighting. In contrast, equal-weighted means are notably higher for the Dimson (DIM), Scholes-Williams (SW), and Frazzini-Pederson betas (FP). These three estimators also exhibit greater dispersion, both cross-sectionally and over time, suggesting that incorporating lagged or lead market betas introduces more cross-sectional variation in estimated betas. Estimators that apply shrinkage or winsorization, such as Blume (BLM) and Vasicek (VCK), generate more stable betas. Realized beta (RB) has the widest range and highest variation. Panel B echoes the patterns from Panel A, showing that most beta estimators are highly correlated, while correlations are notably lower for DIM, SW, FP, and RB.

<TABLE 4 HERE>

Table 4 reports estimates of γ for all adjusted regression-based betas, benchmarked against the 24-month OLS beta (OLS_D^{24}).[5] As a linear function of the OLS beta, Blume beta (BLM) yields an identical γ of 0.346 by construction. Most alternative estimators have significantly higher γ values relative to the OLS benchmark, showing better alignment with the true, unobserved beta.

We observe a clear tier in improvements. Adjustments that reduce extreme noise, such as winsorization and shrinkage (e.g., BSW and VCK) deliver modest improvements in accuracy. Their γ values are only economically slightly above the OLS benchmark. This suggests that trimming outliers helps, but offers limited upside. A second tier includes estimators that put greater weight on recent returns, such as $BSWA$ and $EWLS$. Their γ values show more substantial improvements, suggesting that weighting recent data enhances responsiveness to the time-varying risk exposures for individual stocks.

[5] In the Internet Appendix, we show that results are similar when using the 12-month OLS beta as the benchmark.

The best performing individual estimator is the Frazzini-Pedersen beta FP. Finally, the simple average of all ten estimators (AVG) delivers the strongest performance. Importantly, this average is taken across estimators that each contribute meaningful improvements. The resulting diversification effect reduces idiosyncratic noise and generates an estimate better correlated with the true, unobserved beta. This finding is related to the model averaging literature, which documents improved prediction accuracy by combining estimates from multiple models rather than relying on a single specification (e.g., Avramov, 2002; Rapach et al., 2010).

4.4 Option-implied betas

In addition to the regression-based betas above, we also consider the option-implied betas ($\widetilde{\beta}_{i,t}^{BV,k}$) proposed by Buss and Vilkov (2012). These betas are derived from implied volatilities and correlations from stock and index option prices, providing a forward-looking measure of market beta. We obtain option-implied beta estimates for 862 stocks over 323 months (February 1997 to December 2023), with a total sample size of 100,185 stock-month observations.[6] Following Buss and Vilkov (2012), we consider option-implied betas implied by options with maturities (k) of 30, 91, 182, 273, and 365 days. We also calculate the *average* option-implied beta as the simple equal-weighted average across the five maturities.

<TABLE 5 HERE>

Table 5 provides summary statistics for the option-implied betas across maturities and their simple average. As a benchmark, we also report the 24-month OLS beta for the

[6] Daily option-implied beta estimates are kindly provided by Grigory Vilkov at `https://osf.io/z2486/`.

subsample for which option-implied betas are available. In Panel A, both equal-weighted and value-weighted means are close to one, consistent with the fact that the sample consists mostly of large cap stocks with actively traded options. Option-implied betas show lower cross-sectional and time-series variation than regression-based betas, as they capture market expectations of future risks. In addition, the standard deviations decline with maturity, suggesting that options with longer maturity capture more persistent components of beta. The average beta across maturities, BV_{AVG} balances the short and long-term expectations and delivers an aggregate forward-looking measure of a stock's market risk.

Panel B shows that option-implied betas across maturities are highly correlated. This suggests strong consistency in the expectations extracted from option markets. Correlations with the OLS benchmark are lower, especially for betas from the short-maturity options. This gap reflects differences between backward-looking historical estimates and the forward-looking nature of option-implied measures.

<TABLE 6 HERE>

Table 6 reports the γ estimates for the option-implied betas. All option-based estimators outperform the OLS benchmark, with statistically significant improvement for maturities of 91 days and longer. The beta derived from 182-day options has the highest value of γ. The levels of γ are similar across longer maturities, consistent with the high correlations among these measures. This also suggests that the market's forward-looking expectations of risk are persistent across horizons beyond short-dated options. The average option-implied beta offers no further improvement, likely because the individual estimates share substantial common information.

4.5 Portfolio beta

In the literature on cross-sectional asset pricing, it is common to proxy the true, unobserved beta of an individual stock using the estimated beta of a characteristic-sorted portfolio, to which the stock is assigned. Fama and French (1992) argue that using portfolio-level betas provides more precise estimates; therefore, mitigating the errors-in-variables (EIV) bias in cross-sectional regressions. However, this approach does introduce a tradeoff, since estimating beta at the portfolio level reduces stock-level idiosyncratic information and cross-sectional variation in betas (See, e.g. Jegadeesh et al., 2019, and Ang et al., 2020).

In this section, we examine whether assigning portfolio-level betas to individual stocks improves estimation accuracy compared to the benchmark OLS beta estimated at the stock level. Each month, we sort all stocks into J portfolios based on their 24-month OLS beta estimates from daily returns.[7] We then compute value-weighted portfolio returns and estimate portfolio betas by regressing daily excess portfolio returns on the market excess return over a 24-month rolling window. These portfolio-level betas are assigned back to each stock. Finally, we compute γ using Eq. (6).

<TABLE 7 HERE>

Table 7 reports the γ estimates for portfolio-level betas assigned to individual stocks, based on different choices of J. We compare these portfolio betas against the benchmark OLS beta estimated at the stock level from daily returns. With a small number of portfolios the portfolio betas are significantly less accurate than the stock-level betas. As the number of portfolios, J, increases from 2 to 100, the assigned portfolio-level betas gradually converge

[7] In the Internet Appendix, we repeat this analysis using betas estimated over 12-month windows, leading to very similar results.

to the stock-level OLS benchmark in accuracy. However, none of the portfolio-based betas outperforms the stock-level beta. Our analysis therefore suggests that the practice of grouping stocks into portfolios and assigning portfolio betas back to the individual stocks does not lead to more accurate betas.

4.6 Comparison to realized beta

In addition to our γ measure, we also benchmark the estimated betas against realized betas. This approach, commonly used in the literature (e.g., Hollstein and Prokopczuk, 2016, Cosemans et al., 2016, Welch, 2022), implicitly assumes realized beta as a proxy for the true beta. To implement this comparison, we use the definition in (15) and compute two standard evaluation metrics. First, we evaluate the average cross-sectional correlation between the estimated and realized betas:

$$CORR = T^{-1} \sum_{t} \rho_t^{cs} \left[\widetilde{\beta}_{i,t}, \widetilde{\beta}_{i,t}^{RB} \right], \tag{16}$$

where ρ_t^{cs} refers to the cross-sectional correlation coefficient in month t. Second, we evaluate the average root mean square error (RMSE):

$$RMSE = T^{-1} \sum_{t} \sqrt{n_t^{-1} \sum_{i} \left(\widetilde{\beta}_{i,t} - \widetilde{\beta}_{i,t}^{RB} \right)^2}, \tag{17}$$

which is the time-series average of the square root of the MSE defined in (14).

<TABLE 8 HERE>

Table 8 evaluates the accuracy of each beta estimator using three metrics: CORR, RMSE, and our proposed γ measure. The table reports values and rankings for all estimators across the four beta categories introduced in previous sections.

We find that CORR and RMSE metrics generate highly similar rankings across beta estimators. Within each group of OLS betas, adjusted regression-based betas, and portfolio betas, the rank correlations between CORR and RMSE are 0.93, 0.99, and 0.93, respectively. As discussed in Section 2.2, the similarity between CORR and RMSE suggests that differences in accuracy between beta estimators mostly reflect differences in noise, rather than bias.

In contrast, rankings based on our proposed γ measure diverge substantially, particularly for the adjusted regression-based betas. The rank correlations between γ and CORR is only -0.19 and between γ and RMSE is -0.15. For example, $EWLS_{(84)}$ and $EWLS_{(168)}$ are the top-ranked beta estimators based on CORR and RMSE but drop to mid-tier based on the γ measure. Conversely, DIM and SW betas rank poorly when benchmarked against realized beta comparisons, but rank much better using γ. This illustrates the problem with benchmarking against realized beta. Both Dimson and Scholes-Williams betas are designed to capture lagged adjustment to market returns for thinly traded stocks. Realized beta, which does not account for such lagged adjustment, may therefore be a poor proxy of true beta for such illiquid stocks. This results in Dimson and Scholes-Williams betas performing poorly when benchmarked against realized beta, even when the correlation with true beta is high. Realized beta, while estimated in-sample, is still an estimate of beta, and estimation choices affect the ranking of historical beta estimates. Our approach, on the other hand, does not require any assumptions about the true beta.

We find similar patterns in the portfolio betas, where rankings under CORR and RMSE are consistent but deviate from those under γ. For option-implied betas, the rankings are broadly similar across all measures, though correlations are somewhat lower due to the smaller number of estimators.

Taken together, Table 8 highlights that our γ measure provides a distinct and more robust perspective on the estimation accuracy of market beta estimates. Unlike realized beta, which is itself subject to estimation error and does not account for features such as lagged adjustment or long horizon dynamics, our γ measure evaluates the estimates without requiring a proxy for the true beta.

4.7 Mutual fund betas

In the mutual funds literature, a common practice is to construct predictive alpha, defined as the difference between a fund's realized return and its expected return, where the expected return is based on the fund's historical factor exposures interacted with contemporaneous factor realizations. Predictive alpha is often used as a forward-looking measure of abnormal performance, and its construction assumes that lagged betas serve as reasonable proxies for a fund's expected risk exposures.

This framework is widely applied in studies whether certain fund characteristics, such as R2 (Amihud and Goyenko, 2013), managerial active weight (Doshi et al., 2015) or past returns (e.g., Carhart, 1997; Busse et al., 2010; Bali et al., 2011), are associated with superior future performance. Cederburg et al. (2018) also uses lagged betas as instruments in conditional performance evaluation models to account for time-varying exposures. While the estimation approach varies, all rely on rolling historical return windows to estimate beta.

For example, Amihud and Goyenko (2013) use historical 24 months of monthly return and 6 months of daily return with Dimson adjustment. Carhart (1997) and Bali et al. (2011) use 36-month windows, while Grønborg et al. (2021) extends the horizon to 60 months. Despite the variation in method, horizon and frequency, these studies share a common assumption that the beta estimates using historical return data reliably reflect a fund's conditional exposures at the time of performance evaluation. In this section, we revisit this assumption with our γ framework.

We construct our sample from the CRSP Survivor-Bias-Free Mutual Fund database. We focus on actively managed U.S. equity funds.[8] Fund returns are net of expenses. To address the incubation and back-fill bias (Evans, 2010), we keep only funds that are at least 2 years old and have at least \$15 million in assets. Given that CRSP's mutual fund return data begin in September 1998, our final sample starts in September 1999, runs until December 2023, and comprises 3,098 unique funds.

Following the literature, we consider a set of standard estimators based on both monthly and daily fund returns. Using monthly returns, we estimate OLS and Vasicek-adjusted betas (e.g., Irvine et al., 2024) over 36 and 60-month rolling windows. Using daily returns, we compute OLS, Vasicek-adjusted, and Dimson-adjusted betas over 6 and 12-month windows. In addition, we include exponentially weighted least squares betas over a 12-month horizon ($EWLS_{(84)}$ and $EWLS_{(168)}$), similar to our stock analysis above. These estimates capture

[8]To construct the sample, we remove funds that was ever listed as an ETF/Index funds/Variable Annuity/Target allocation/Municipal bond funds. We drop any fund that with less than 75% or more than 105% in equities. U.S. equity funds are defined as those with Lipper codes "EIEI", "LCCE", "LCGE", "LCVE", "MCCE", "MCGE", "MCVE", "MLCE", "MLGE", "MLVE", "SCCE", "SCGE", "SCVE; Strategic Insight Codes "AGG", "GMC", "GRI", "GRO", "ING", "SCG"; Wiesenberger codes "G", "G-I", "GCI", "IEQ", "LTG", "MCG", "SCG".We also remove target date funds and other international, emerging market, balanced, or similar funds (etc.), based on the fund name. The list of key words is available upon request.

a range of assumptions regarding horizon length, responsiveness to new information, and shrinkage to cross-sectional mean. For each fund-month observation, we also compute the simple average across these twelve beta estimates, which we refer to as the average beta.

<TABLE 9 HERE>

Table 9 presents the γ estimates for various beta estimators, using the 60-month OLS beta estimated from monthly returns as the benchmark. We observe several patterns. First, Vasicek-adjusted and EWLS betas exhibit slightly higher γ values compared to their OLS counterparts, but the improvements are modest and statistically insignificant. In contrast, Dimson (1979) beta outperforms all other estimators, with significantly higher γ values. Although Cremers et al. (2013) argue that the staleness in fund returns is likely to be similar to benchmark index returns, our results suggest that explicitly correcting for staleness improves beta accuracy. Second, we observe a frequency-specific trade-off in the choice of estimation window. For monthly returns, the 60-month beta outperforms the 36-month beta, underscoring the importance of sufficient observations for accurate estimation. In contrast, for daily returns, where data are abundant, the 6-month window performs better than the 12-month alternative. This highlights the benefits of responsiveness to recent information in the actively managed mutual fund setting.

5 Conclusion

In this paper, we introduce a novel framework to evaluate the accuracy of market beta estimates. This framework addresses a long-standing challenge of comparing beta estimators when the true betas are unobserved. Our proposed measure, γ, requires only minimal assumptions about the asset pricing model and is applicable to any beta estimates regardless of the underlying methodology. By being proportional to the cross-sectional correlation between the estimated and true market betas, γ offers a consistent and model-agnostic tool to rank beta estimators by accuracy.

We validate our measure through a comprehensive set of Monte Carlo simulations that span a wide range of asset pricing models, including settings with priced or unpriced market risk and potential correlation with other priced or unpriced factors. We show that γ closely tracks the correlation between estimated and true betas and exhibits high power in detecting even modest differences in estimation accuracy.

Applying the framework to U.S. stocks, we find that betas estimated from daily returns are consistently more accurate than their monthly counterparts. Adjustments such as shrinkage, exponential weighting, and corrections for nonsynchronous trading systematically improve estimation accuracy. Notably, simply averaging these improved estimators delivers the strongest performance, highlighting the value of diversifying estimation noise through model averaging. We further examine option-implied betas in a smaller sample.

Given the widespread use of portfolio-assigned betas to mitigate errors-in-variables (EIV) bias, we evaluate whether this approach improves estimation accuracy relative to stock-level betas. Across all specifications, we find portfolio-level proxies yield lower γ values than their stock-level counterparts, particularly when portfolios contain a large number of stocks. While increasing the number of portfolios reduces the gap, we find no setting in which portfolio-based betas outperform direct stock-level estimates. Our evidence is consistent with recent literature (e.g., Jegadeesh et al., 2019) suggesting that portfolio formation conceals relevant cross-sectional variation.

A key finding of our paper is that γ yields a ranking of beta estimators that often diverges from traditional approaches based on realized, in-sample betas. This discrepancy arises because conventional measures implicitly assume that realized betas are accurate proxies for the true betas, though they may be subject to estimation error and model misspecification. In contrast, γ avoids this assumption and offers a more robust assessment of estimation accuracy.

Finally, we use U.S. actively managed equity mutual funds as a practical setting to evaluate which beta estimators best capture funds' expected market exposure for predictive alpha. Among the alternatives, we find that Dimson-adjusted beta using daily return data achieves the highest γ. We also observe a clear trade-off between data frequency and estimation window, as longer horizons improve accuracy with monthly returns, whereas shorter and more responsive windows perform better with daily data.

While our framework highlights meaningful differences in the accuracy of beta estimators, we do not prescribe a single "best" estimator, as the ranking may vary across samples, contexts, and market conditions. Instead, we recommend that researchers and practitioners adopt our framework for evaluating and comparing beta estimators. Our γ measure provides a simple tool for comparing beta estimators in asset pricing, corporate finance, and investment applications.

Appendix A Unbiasedness of γ

Define $y_t \equiv \dfrac{Cov_t^{cs}\left[\widetilde{\beta}_{i,t}, R_{i,t}\right]}{\sigma_t^{cs}\left[\widetilde{\beta}_{i,t}\right]}$, $\gamma_t \equiv \dfrac{Cov_t^{cs}\left[\widetilde{\beta}_{i,t}, \beta_{i,t}\right]}{\sigma_t^{cs}\left[\widetilde{\beta}_{i,t}\right]}$, and $\zeta_t \equiv \dfrac{Cov_t^{cs}\left[\widetilde{\beta}_{i,t}, e_{i,t}\right]}{\sigma_t^{cs}\left[\widetilde{\beta}_{i,t}\right]}$,
where $e_{i,t} \equiv R_{i,t} - \beta_{i,t} R_{m,t}$, such that equation (4) reduces to

$$y_t = \gamma_t R_{m,t} + \zeta_t, \tag{A.1}$$

and the time-series regression (6) becomes:

$$y_t = \delta + \gamma R_{m,t} + \xi_t. \tag{A.2}$$

For the OLS estimator of γ to identify the mean of γ_t, we need to first assume that γ_t is a stationary process with mean $\bar{\gamma}$, such that $\gamma_t = \bar{\gamma} + \psi_t$, where ψ_t is a mean-zero process. Substituting (A.1) and $\gamma_t = \bar{\gamma} + \psi_t$ into the OLS estimator of γ gives:

$$E\left[\widehat{\gamma}\right] = \frac{Cov^{ts}\left[y_t, R_{m,t}\right]}{Var^{ts}\left[R_{m,t}\right]} = \frac{Cov^{ts}\left[\left(\bar{\gamma} + \psi_t\right) R_{m,t} + \zeta_t, R_{m,t}\right]}{Var^{ts}\left[R_{m,t}\right]} \tag{A.3}$$

$$= \bar{\gamma} + \frac{Cov^{ts}\left[\psi_t R_{m,t}, R_{m,t}\right]}{Var^{ts}\left[R_{m,t}\right]} + \frac{Cov^{ts}\left[\zeta_t, R_{m,t}\right]}{Var^{ts}\left[R_{m,t}\right]}, \tag{A.4}$$

where $Var^{ts}[.]$ and $Cov^{ts}[.]$ refer to the time-series variance and covariance, respectively. This means that $\widehat{\gamma}$ is an unbiased estimator of $\bar{\gamma}$ ($E\left[\widehat{\gamma}\right] = \bar{\gamma}$) under three conditions. First, the cross-sectional covariance $Cov_t^{cs}\left[\widetilde{\beta}_{i,t}, e_{i,t}\right]$ is uncorrelated with $R_{m,t}$. Second, the cross-sectional standard deviation of estimated beta $\sigma_t^{cs}\left[\widetilde{\beta}_{i,t}\right]$ is uncorrelated with $R_{m,t}$. Third, ψ_t (measuring time-variation in γ_t) is uncorrelated with $R_{m,t}$. The analysis below shows that the first condition is equivalent to the testable condition that estimated betas are uncorrelated with $R_{m,t}$.

For ease of exposition, the number of stocks N is assumed constant over time and market weights of each stock are normalized to $\frac{1}{N}$, such that $R_{m,t} = N^{-1}\sum_{i=1}^{N} R_{i,t}$. Using the decomposition (1), we can show that the cross-sectional mean of $e_{i,t}$ is therefore zero in each period:[9]

$$N^{-1}\sum_{i=1}^{N} e_{i,t} = \underbrace{N^{-1}\sum_{i=1}^{N} R_{i,t}}_{R_{m,t}} - R_{m,t}\underbrace{N^{-1}\sum_{i=1}^{N} \beta_{i,t}}_{1} = 0$$

This gives:

$$\begin{aligned}
Cov_t^{cs}\left[\widetilde{\beta}_{i,t}, e_{i,t}\right] &= N^{-1}\sum_{i=1}^{N} \widetilde{\beta}_{i,t} e_{i,t} - N^{-1}\sum_{i=1}^{N} \widetilde{\beta}_{i,t}\underbrace{N^{-1}\sum_{i=1}^{N} e_{i,t}}_{0} \\
&= N^{-1}\sum_{i=1}^{N} \widetilde{\beta}_{i,t} e_{i,t}. \\
Cov^{ts}\left[R_{m,t}, Cov_t^{cs}\left[\widetilde{\beta}_{i,t}, e_{i,t}\right]\right] &= Cov^{ts}\left[R_{m,t}, N^{-1}\sum_{i=1}^{N} \widetilde{\beta}_{i,t} e_{i,t}\right].
\end{aligned}$$

Without loss of generalization, we can assume that $R_{m,t}$ has a mean of zero. If $Cov_t^{cs}\left[\widetilde{\beta}_{i,t}, e_{i,t}\right]$ is uncorrelated with $R_{m,t}$, it is also uncorrelated with $R_{m,t} + c$.

$$\begin{aligned}
Cov^{ts}\left[R_{m,t}, N^{-1}\sum_{i=1}^{N} \widetilde{\beta}_{i,t} e_{i,t}\right] &= T^{-1}\sum_{t=1}^{T} R_{m,t} N^{-1}\sum_{i=1}^{N} \widetilde{\beta}_{i,t} e_{i,t} \\
&= N^{-1}T^{-1}\sum_{t=1}^{T}\sum_{i=1}^{N} \widetilde{\beta}_{i,t} e_{i,t} R_{m,t} \\
&= N^{-1}\sum_{i=1}^{N} T^{-1}\sum_{t=1}^{T} \widetilde{\beta}_{i,t} e_{i,t} R_{m,t}.
\end{aligned}$$

The last part of this equation, the time-series mean of $\widetilde{\beta}_{i,t} e_{i,t} R_{m,t}$ can be expressed as:

[9] Note: $e_{i,t}$ can have a non-zero *time-series* mean for a given stock i. For example, a positive alpha stock will have a positive time-series mean of $e_{i,t}$.

$$
\begin{aligned}
T^{-1}\sum_{t=1}^{T}\widetilde{\beta}_{i,t}e_{i,t}R_{m,t} &= T^{-1}\sum_{t=1}^{T}\widetilde{\beta}_{i,t}T^{-1}\sum_{t=1}^{T}e_{i,t}\underbrace{T^{-1}\sum_{t=1}^{T}R_{m,t}}_{0} \\
&\quad +\underbrace{Cov_i^{ts}\left[e_{i,t},R_{m,t}\right]}_{0}T^{-1}\sum_{t=1}^{T}\widetilde{\beta}_{i,t} \\
&\quad +Cov_i^{ts}\left[\widetilde{\beta}_{i,t},e_{i,t}\right]\underbrace{T^{-1}\sum_{t=1}^{T}R_{m,t}}_{0} \\
&\quad +Cov_i^{ts}\left[\widetilde{\beta}_{i,t},R_{m,t}\right]T^{-1}\sum_{t=1}^{T}e_{i,t} \\
&= Cov_i^{ts}\left[\widetilde{\beta}_{i,t},R_{m,t}\right]T^{-1}\sum_{t=1}^{T}e_{i,t}
\end{aligned}
$$

Therefore, the covariance of interest between the market return and $Cov_t^{cs}\left[\widetilde{\beta}_{i,t},e_{i,t}\right]$ thus becomes:

$$
\begin{aligned}
Cov^{ts}\left[R_{m,t},N^{-1}\sum_{i=1}^{N}\widetilde{\beta}_{i,t}e_{i,t}\right] &= N^{-1}\sum_{i=1}^{N}T^{-1}\sum_{t=1}^{T}\widetilde{\beta}_{i,t}e_{i,t}R_{m,t} \\
&= N^{-1}\sum_{i=1}^{N}Cov_i^{ts}\left[\widetilde{\beta}_{i,t},R_{m,t}\right]T^{-1}\sum_{t=1}^{T}e_{i,t}
\end{aligned}
\tag{A.5}
$$

This covariance is equal to zero if $Cov_i^{ts}\left[\widetilde{\beta}_{i,t},R_{m,t}\right]$ is zero, which is a testable assumption. For each stock i, and for each of our regression-based betas, we estimate the following predictive time-series regression:

$$
R_{m,t}=\phi_0+\phi_1\widetilde{\beta}_{i,t}+\epsilon_{i,t}, \tag{A.6}
$$

and test the hypothesis $H_0:\phi_1=0$. Under this null, the covariance in (A.5) is equal to zero, which is the first condition for γ to be unbiased. Panel A of Table A.1 reports regression frequencies for t-tests at the 10%, 5% , and 1% significance level, for each of the regression-based betas.

<TABLE A.1 HERE>

The results show that the rejection frequencies are close to the nominal significance levels, suggesting indeed that the null holds and estimated stock-level betas are not correlated with market-returns. The result that *stock-level* betas $\widetilde{\beta}_{i,t}$, estimated using data prior to period t, do not hold predictive power for the market return at time t, is not surprising since market returns are difficult to predict using lagged information.

To test the second condition for unbiasedness of γ, we also estimate the following predictive time-series regression for each of our regression-based betas:

$$R_{m,t} = \psi_0 + \psi_1 \sigma_t^{cs} \left[\widetilde{\beta}_{i,t} \right] + u_t. \tag{A.7}$$

Under the null $H_0 : \psi_1 = 0$, the cross-sectional standard deviation $\sigma_t^{cs} \left[\widetilde{\beta}_{i,t} \right]$ is uncorrelated with $R_{m,t}$. Estimates and standard errors of ψ_1 are reported in Panel B of Table A.1. As expected, $\sigma_t^{cs} \left[\widetilde{\beta}_{i,t} \right]$, which is based on lagged information, is not significantly correlated with the market return $R_{m,t}$.

The lack of predictability of market returns by lagged information (including market betas estimated using historical returns) also suggests that it in unlikely that time-variation in the correlation between estimated and true betas is correlated with $R_{m,t}$, which is outside the estimation window of the betas. This in turn implies that also the third condition implied by Eq. (A.3) (ψ_t is uncorrelated with $R_{m,t}$) is satisfied and that the OLS regression (6) provides an unbiased and consistent estimate of γ.

References

Amihud, Y. and Goyenko, R. (2013). Mutual fund's r^2 as predictor of performance. *The Review of Financial Studies*, 26(3):667–694.

Andersen, T. G., Bollerslev, T., Diebold, F. X., and Wu, G. (2006). Realized beta: Persistence and predictability. In *Econometric Analysis of Financial and Economic Time Series*, pages 1–39. Emerald Group Publishing Limited.

Ang, A., Liu, J., and Schwarz, K. (2020). Using stocks or portfolios in tests of factor models. *Journal of Financial and Quantitative Analysis*, 55(3):709–750.

Avramov, D. (2002). Stock return predictability and model uncertainty. *Journal of Financial Economics*, 64(3):423–458.

Bali, T. G., Brown, S. J., and Caglayan, M. O. (2011). Do hedge funds' exposures to risk factors predict their future returns? *Journal of Financial Economics*, 101(1):36–68.

Blume, M. E. (1971). On the assessment of risk. *Journal of Finance*, 26(1):1–10.

Boons, M. (2016). State variables, macroeconomic activity, and the cross section of individual stocks. *Journal of Financial Economics*, 119(3):489–511.

Buss, A. and Vilkov, G. (2012). Measuring equity risk with option-implied correlations. *The Review of Financial Studies*, 25(10):3113–3140.

Busse, J. A., Goyal, A., and Wahal, S. (2010). Performance and persistence in institutional investment management. *The Journal of Finance*, 65(2):765–790.

Carhart, M. M. (1997). On persistence in mutual fund performance. *The Journal of Finance*, 52(1):57–82.

Cederburg, S., O'Doherty, M. S., Savin, N. E., and Tiwari, A. (2018). Conditional benchmarks and predictors of mutual fund performance. *Critical Finance Review*, 7(2):331–372.

Cosemans, M., Frehen, R., Schotman, P. C., and Bauer, R. (2016). Estimating security betas using prior information based on firm fundamentals. *Review of Financial Studies*, 29(4):1072–1112.

Cremers, M., Petajisto, A., and Zitzewitz, E. (2013). Should benchmark indices have alpha? revisiting performance evaluation. *Critical Finance Review*, 2(1):1–48.

Da, Z., Guo, R.-J., and Jagannathan, R. (2012). Capm for estimating the cost of equity capital: Interpreting the empirical evidence. *Journal of Financial Economics*, 103(1):204–220.

Dimson, E. (1979). Risk measurement when shares are subject to infrequent trading. *Journal of Financial Economics*, 7(2):197–226.

Doshi, H., Elkamhi, R., and Simutin, M. (2015). Managerial activeness and mutual fund performance. *The Review of Asset Pricing Studies*, 5(2):156–184.

Décaire, P. H., Sosyura, D., and Wittry, M. D. (2024). Resolving estimation ambiguity. Fisher College of Business Working Paper No. 2024-03-019.

Eaton, G. W., Guo, F., Liu, T., and Tu, D. (2025). The cost of equity: Evidence from investment banking valuations. *Journal of Financial and Quantitative Analysis*, 60(1):1–39.

Evans, R. B. (2010). Mutual fund incubation. *The Journal of Finance*, 65(4):1581–1611.

Fama, E. and French, K. (1992). The cross-section of expected stock returns. *Journal of Finance*, 47(2):427–465.

Fama, E. F. and French, K. R. (1997). Industry costs of equity. *Journal of Financial Economics*, 43(2):153–193.

Fama, E. F. and MacBeth, J. D. (1973). Risk, return, and equilibrium: Empirical tests. *Journal of Political Economy*, 81(3):607–636.

Frazzini, A. and Pedersen, L. H. (2014). Betting against beta. *Journal of Financial Economics*, 111(1):1–25.

Graham, J. R. and Harvey, C. R. (2001). The theory and practice of corporate finance: Evidence from the field. *Journal of Financial Economics*, 60(2-3):187–243.

Grønborg, N. S., Lunde, A., Timmermann, A., and Wermers, R. (2021). Picking funds with confidence. *Journal of Financial Economics*, 139(1):1–28.

Guo, H., Wu, C., and Yu, Y. (2017). Time-varying beta and the value premium. *Journal of Financial and Quantitative Analysis*, 52(4):1551–1576.

Hollstein, F. and Prokopczuk, M. (2016). Estimating beta. *Journal of Financial and Quantitative Analysis*, 51(4):1437–1466.

Hollstein, F., Prokopczuk, M., and Simen, C. W. (2020). Beta uncertainty. *Journal of Banking & Finance*, 116:105834.

Irvine, P., Kim, J. H., and Ren, J. (2024). The beta anomaly and mutual fund performance. *Management Science*, 70(1):143–163.

Jegadeesh, N., Noh, J., Pukthuanthong, K., Roll, R., and Wang, J. (2019). Empirical tests of asset pricing models with individual assets: Resolving the errors-in-variables bias in risk premium estimation. *Journal of Financial Economics*, 133(2):273–298.

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

Markowitz, H. (1952). The utility of wealth. *Journal of Political Economy*, 60(2):151–158.

Patton, A. J. (2011). Volatility forecast comparison using imperfect volatility proxies. *Journal of Econometrics*, 160(1):246–256.

Rapach, D. E., Strauss, J. K., and Zhou, G. (2010). Out-of-sample equity premium prediction: Combination forecasts and links to the real economy. *The Review of Financial Studies*, 23(2):821–862.

Scholes, M. and Williams, J. (1977). Estimating betas from nonsynchronous data. *Journal of Financial Economics*, 5(3):309–327.

Sharpe, W. F. (1964). Capital asset prices: A theory of market equilibrium under conditions of risk. *Journal of Finance*, 19(3):425–442.

Vasicek, O. A. (1973). A note on using cross-sectional information in bayesian estimation of security betas. *Journal of Finance*, 28(5):1233–1239.

Welch, I. (2022). Simply better market betas. *Critical Finance Review*, 11(1):37–64.

Welch, I. and Goyal, A. (2008). A comprehensive look at the empirical performance of equity premium prediction. *The Review of Financial Studies*, 21(4):1455–1508.

Figure 1: Simulations. This figure plots the correlation ρ between true and estimated beta (x-axis) against estimated γ (y-axis), for $R = 100$ panels of pseudo-estimated betas.

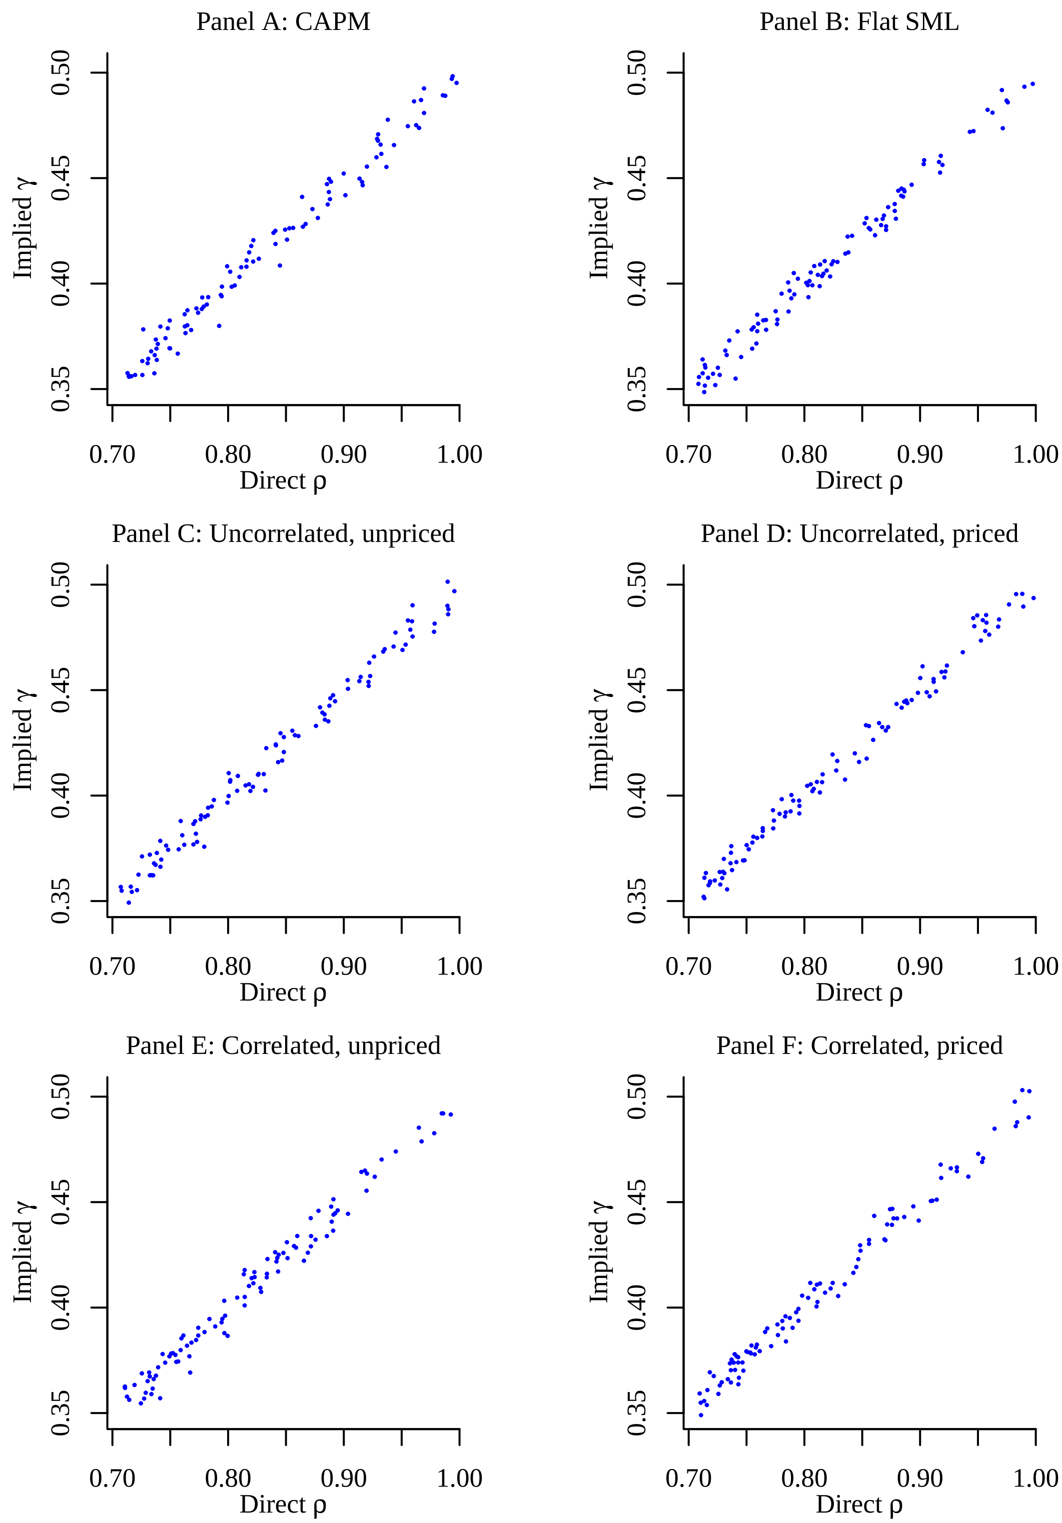

Figure 2: Power curves. This figure plots 5% (red line) and 1% (blue line) rejection frequencies of the null-hypothesis $H_0 : \Delta\gamma = 0$ (y-axis) against the true difference in correlation with true beta between two panels of pseudo-estimated betas (x-axis).

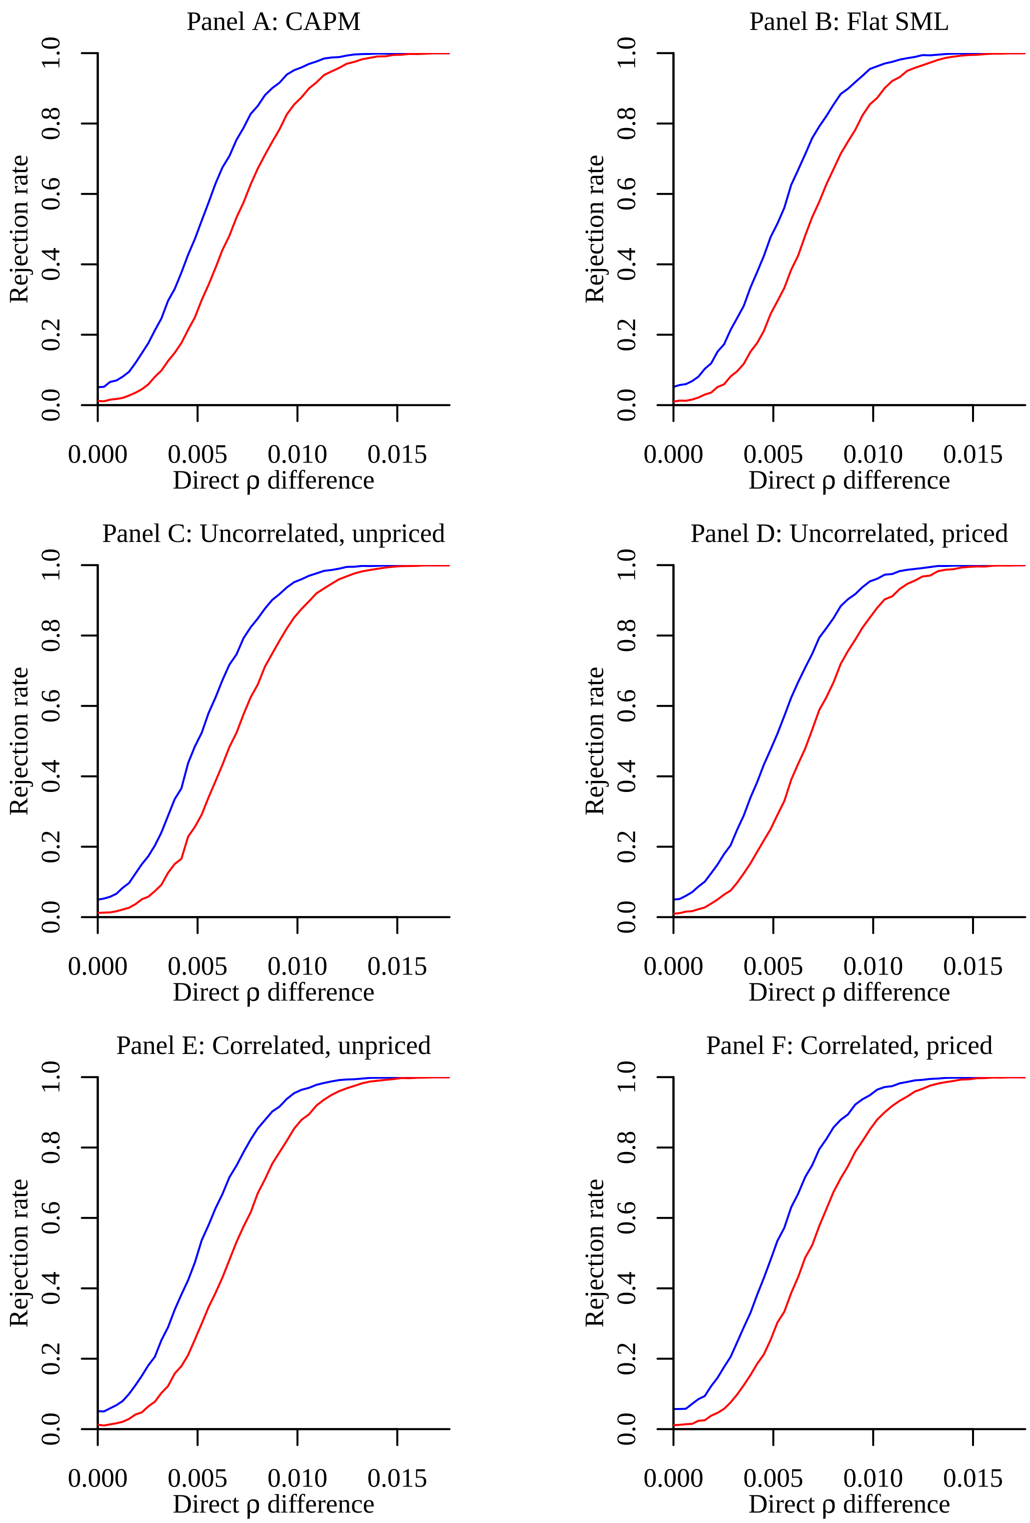

Table 1: Summary statistics - OLS betas

This table presents summary statistics for betas estimated by OLS. OLS_D^K refers to betas estimated by rolling-window OLS using K months of daily return observations. OLS_M^K refers to betas estimated by rolling-window OLS using K monthly return observations. See Section 4.2 for details. The number of firm-month observations is $1,974,590$, spanning $T = 552$ months (1978:1-2023:12) and $N = 16,342$ unique firms. The row $Mean^{EW}$ reports for each beta the time-series mean of the equal-weighted cross-sectional average beta in each period. $Mean^{VW}$ is time-series mean of the vale-weighted cross-sectional average beta in each period. $Std.Dev.^{CS}$ is the time-series average of the cross-sectional standard deviation. $Std.Dev.^{TS}$ is the cross-sectional average of the firm-level time-series standard deviation of beta. The final five rows report selected quantiles from the pooled panel of betas. Panel B reports correlations between the various betas. Each cell represents the time-series average of the cross-sectional correlation coefficient.

Panel A: Summary statistics

	OLS_D^{12}	OLS_D^{6}	OLS_D^{24}	OLS_D^{36}	OLS_D^{60}	OLS_M^{12}	OLS_M^{24}	OLS_M^{36}	OLS_M^{60}
Mean^{EW}	0.819	0.816	0.817	0.815	0.806	1.074	1.091	1.103	1.103
Mean^{VW}	1.023	1.021	1.025	1.026	1.024	0.998	0.997	0.999	1.003
Std.Dev.^{CS}	0.576	0.670	0.516	0.490	0.461	1.370	0.971	0.832	0.708
Std.Dev.^{TS}	0.371	0.519	0.265	0.218	0.165	1.146	0.674	0.500	0.335
q5	-0.049	-0.171	0.021	0.049	0.078	-0.762	-0.226	-0.044	0.107
q25	0.348	0.314	0.373	0.386	0.398	0.332	0.488	0.549	0.606
Median	0.755	0.751	0.757	0.756	0.745	0.952	0.985	1.003	1.017
q75	1.182	1.210	1.158	1.145	1.119	1.656	1.543	1.509	1.465
q95	1.862	1.988	1.758	1.706	1.642	3.309	2.771	2.568	2.362

Panel B: Correlations

	OLS_D^{12}	OLS_D^{6}	OLS_D^{24}	OLS_D^{36}	OLS_D^{60}	OLS_M^{12}	OLS_M^{24}	OLS_M^{36}	OLS_M^{60}
OLS_D^{12}	1.000								
OLS_D^{6}	0.866	1.000							
OLS_D^{24}	0.891	0.761	1.000						
OLS_D^{36}	0.829	0.708	0.946	1.000					
OLS_D^{60}	0.758	0.647	0.869	0.934	1.000				
OLS_M^{12}	0.395	0.351	0.366	0.347	0.323	1.000			
OLS_M^{24}	0.456	0.394	0.499	0.478	0.447	0.744	1.000		
OLS_M^{36}	0.474	0.408	0.533	0.555	0.526	0.634	0.859	1.000	
OLS_M^{60}	0.478	0.410	0.545	0.582	0.616	0.530	0.723	0.846	1.000

Table 2: γ estimates - OLS betas

This table reports γ, estimated by regressing monthly observations of $\frac{Cov_t^{cs}[\widetilde{\beta}_{i,t},R_{i,t}]}{\sigma_t^{cs}[\widetilde{\beta}_{i,t}]}$ on the monthly excess market return (See Eq. (6)), for various panels of OLS beta estimates $\widetilde{\beta}_{i,t}$. γ is proportional to the correlation between the empirically estimated betas and the unobservable true beta. The first two rows report OLS estimates of γ and corresponding Newey-West standard errors (SE[γ]), for each of the different beta estimators. $\Delta\gamma$ and SE[$\Delta\gamma$] report the difference between γ of each beta and the corresponding subsample of daily OLS beta, along with their Newey-West standard errors (See Eq. (7)).

	OLS_D^{12}	OLS_D^{6}	OLS_D^{24}	OLS_D^{36}	OLS_D^{60}	OLS_M^{12}	OLS_M^{24}	OLS_M^{36}	OLS_M^{60}
γ	0.340	0.311	0.346	0.346	0.342	0.236	0.286	0.311	0.327
SE[γ]	0.028	0.027	0.028	0.027	0.025	0.029	0.030	0.029	0.028
$\Delta\gamma$	.	-0.028	0.007	0.006	0.003	-0.104	-0.054	-0.029	-0.013
SE[$\Delta\gamma$]	.	0.004	0.005	0.007	0.009	0.014	0.013	0.014	0.013

Table 3: Summary statistics - Adjusted OLS betas

This table presents summary statistics for adjusted regression-based betas. See Section 4.3 for details. The number of firm-month observations is $1,974,590$, spanning $T = 552$ months (1978:1-2023:12) and $N = 16,342$ unique firms. The row $Mean^{EW}$ reports for each beta the time-series mean of the equal-weighted cross-sectional average beta in each period. $Mean^{VW}$ is time-series mean of the vale-weighted cross-sectional average beta in each period. $Std.Dev.^{CS}$ is the time-series average of the cross-sectional standard deviation. $Std.Dev.^{TS}$ is the cross-sectional average of the firm-level time-series standard deviation of beta. The final five rows report selected quantiles from the pooled panel of betas. Panel B reports correlations between the various betas. Each cell represents the time-series average of the cross-sectional correlation coefficient.

Panel A: Summary statistics

	OLS_D^{24}	BLM	BSW	$BSWA$	VCK	$EWLS_{(84)}$	$EWLS_{(168)}$	DIM	SW	FP	AVG	RB
Mean^{EW}	0.817	0.878	0.820	0.820	0.818	0.821	0.820	1.011	0.914	0.998	0.872	0.806
Mean^{VW}	1.025	1.017	1.007	1.006	1.015	1.005	1.006	1.003	1.012	1.080	1.018	1.016
Std.Dev.^{CS}	0.516	0.344	0.410	0.416	0.452	0.431	0.417	0.636	0.592	0.573	0.441	1.223
Std.Dev.^{TS}	0.265	0.177	0.190	0.208	0.220	0.243	0.209	0.388	0.333	0.328	0.212	1.296
q5	0.021	0.347	0.142	0.132	0.089	0.110	0.132	0.069	0.030	0.149	0.160	-0.935
q25	0.373	0.582	0.459	0.454	0.413	0.442	0.454	0.552	0.453	0.527	0.492	0.143
Median	0.757	0.838	0.771	0.771	0.753	0.770	0.771	0.945	0.840	0.901	0.822	0.734
q75	1.158	1.105	1.107	1.112	1.124	1.122	1.112	1.365	1.256	1.326	1.166	1.368
q95	1.758	1.505	1.578	1.592	1.668	1.622	1.593	2.145	1.991	2.108	1.680	2.693

Panel B: Correlations

	OLS_D^{24}	BLM	BSW	$BSWA$	VCK	$EWLS_{(84)}$	$EWLS_{(168)}$	DIM	SW	FP	AVG	RB
OLS_D^{24}	1.000											
BLM	1.000	1.000										
BSW	0.957	0.957	1.000									
$BSWA$	0.938	0.938	0.979	1.000								
VCK	0.984	0.984	0.978	0.958	1.000							
$EWLS_{(84)}$	0.890	0.890	0.929	0.983	0.910	1.000						
$EWLS_{(168)}$	0.936	0.936	0.977	1.000	0.956	0.984	1.000					
DIM	0.747	0.747	0.717	0.706	0.738	0.673	0.705	1.000				
SW	0.849	0.849	0.823	0.810	0.848	0.773	0.809	0.797	1.000			
FP	0.756	0.756	0.761	0.754	0.763	0.726	0.753	0.677	0.755	1.000		
AVG	0.967	0.967	0.965	0.962	0.972	0.928	0.961	0.826	0.905	0.839	1.000	
RB	0.362	0.362	0.370	0.384	0.369	0.389	0.385	0.279	0.327	0.313	0.376	1.000

Table 4: γ estimates - Adjusted OLS betas

This table reports γ, estimated by regressing monthly observations of $\frac{Cov_t^{cs}[\widetilde{\beta}_{i,t}, R_{i,t}]}{\sigma_t^{cs}[\widetilde{\beta}_{i,t}]}$ on the monthly excess market return (See Eq. (6)), for various panels of adjusted regression-based beta estimates $\widetilde{\beta}_{i,t}$. γ is proportional to the correlation between the empirically estimated betas and the unobservable true beta. The first two rows report OLS estimates of γ and corresponding Newey-West standard errors (SE[γ]), for each of the different beta estimators. $\Delta\gamma$ and SE[$\Delta\gamma$] report the difference between γ of each beta and the baseline daily OLS beta, along with their Newey-West standard errors (See Eq. (7)).

	OLS_D^{24}	BLM	BSW	$BSWA$	VCK	$EWLS_{(84)}$	$EWLS_{(168)}$	DIM	SW	FP	AVG
γ	0.346	0.346	0.354	0.362	0.351	0.360	0.362	0.364	0.362	0.368	0.386
SE[γ]	0.028	0.028	0.028	0.029	0.027	0.029	0.029	0.033	0.030	0.031	0.031
$\Delta\gamma$	.	0.000	0.008	0.015	0.005	0.014	0.015	0.018	0.015	0.022	0.039
SE[$\Delta\gamma$]	.	.	0.003	0.005	0.002	0.007	0.005	0.013	0.007	0.011	0.004

Table 5: Summary statistics - Option-implied betas

This table presents summary statistics for option-implied betas. See Section 4.4 for details. The number of firm-month observations is 100, 185 firm-month observations, spanning $T = 323$ months (1997:2 to 2023:12) and $N = 862$ unique firms. The row $Mean^{EW}$ reports for each beta the time-series mean of the equal-weighted cross-sectional average beta in each period. $Mean^{VW}$ is time-series mean of the vale-weighted cross-sectional average beta in each period. $Std.Dev.^{CS}$ is the time-series average of the cross-sectional standard deviation. $Std.Dev.^{TS}$ is the cross-sectional average of the firm-level time-series standard deviation of beta. The final five rows report selected quantiles from the pooled panel of betas. Panel B reports correlations between the various betas. Each cell represents the time-series average of the cross-sectional correlation coefficient.

Panel A: Summary statistics							
	OLS_D^{24}	BV_{30}	BV_{91}	BV_{182}	BV_{273}	BV_{365}	BV_{AVG}
Mean^{EW}	1.083	1.062	1.072	1.069	1.053	1.047	1.061
Mean^{VW}	1.026	0.999	1.006	1.004	0.997	0.997	1.000
Std.Dev.^{CS}	0.422	0.372	0.337	0.325	0.320	0.312	0.327
Std.Dev.^{TS}	0.196	0.266	0.206	0.190	0.185	0.180	0.197
q5	0.462	0.489	0.555	0.569	0.562	0.564	0.554
q25	0.795	0.810	0.842	0.848	0.839	0.838	0.839
Median	1.043	1.038	1.044	1.043	1.032	1.028	1.039
q75	1.319	1.292	1.277	1.264	1.249	1.239	1.262
q95	1.858	1.760	1.694	1.670	1.651	1.623	1.665
Panel B: Correlations							
	OLS_D^{24}	BV_{30}	BV_{91}	BV_{182}	BV_{273}	BV_{365}	BV_{AVG}
OLS_D^{24}	1.000						
BV_{30}	0.792	1.000					
BV_{91}	0.849	0.942	1.000				
BV_{182}	0.860	0.913	0.987	1.000			
BV_{273}	0.866	0.901	0.977	0.995	1.000		
BV_{365}	0.868	0.892	0.968	0.987	0.995	1.000	
BV_{AVG}	0.860	0.952	0.993	0.992	0.989	0.983	1.000

Table 6: γ estimates - Option-implied betas

This table reports γ, estimated by regressing monthly observations of $\frac{Cov_t^{cs}[\widetilde{\beta}_{i,t},R_{i,t}]}{\sigma_t^{cs}[\widetilde{\beta}_{i,t}]}$ on the monthly excess market return (See Eq. (6)), for various panels of option-implied betas $\widetilde{\beta}_{i,t}$. γ is proportional to the correlation between the empirically estimated betas and the unobservable true beta. The first two rows report OLS estimates of γ and corresponding Newey-West standard errors (SE[γ]), for each of the different beta estimators. $\Delta\gamma$ and SE[$\Delta\gamma$] report the difference between γ of each beta and the baseline daily OLS beta, along with their Newey-West standard errors (See Eq. (7)).

	OLS_D^{24}	BV_{30}	BV_{91}	BV_{182}	BV_{273}	BV_{365}	BV_{AVG}
γ	0.435	0.443	0.463	0.467	0.464	0.461	0.465
SE[γ]	0.043	0.044	0.043	0.043	0.042	0.042	0.043
$\Delta\gamma$	.	0.007	0.027	0.032	0.029	0.026	0.030
SE[$\Delta\gamma$]	.	0.014	0.012	0.012	0.012	0.012	0.012

Table 7: γ estimates - Portfolio betas

This table reports γ for portfolio betas that are assigned to individual stocks. $\widetilde{\beta}_{i,t}^{OLS}$ is the beta of firm i in month t, estimated using 24 months of daily firm-level returns and market returns. $\widetilde{\beta}_{i,t}^{PJ}$ is the beta of a portfolio assigned to firm i in month t. Stocks are each month t sorted into J portfolios based on $\widetilde{\beta}_{i,t}^{OLS}$. Value-weighted daily excess return on each portfolio over the past 24 months are regressed on the daily excess market return to obtain the portfolio beta. The portfolio beta is subsequently assigned to all stocks in the portfolio at time t. Results are reported for sorts into J = 2, 5, 10, 25, 50, 75, and 100 portfolios. The first two rows report OLS estimates of γ and corresponding Newey-West standard errors (SE[γ]), for each of the different betas. $\Delta\gamma$ and SE[$\Delta\gamma$] report the difference between γ of each beta and the baseline daily OLS beta, along with their Newey-West standard errors (See Eq. (7)).

	OLS_D^{24}	$P2$	$P5$	$P10$	$P25$	$P50$	$P75$	$P100$
γ	0.347	0.273	0.330	0.342	0.346	0.347	0.347	0.346
SE[γ]	0.028	0.019	0.026	0.027	0.027	0.028	0.028	0.028
$\Delta\gamma$	.	-0.074	-0.017	-0.005	-0.001	0.000	-0.000	-0.001
SE[$\Delta\gamma$]	.	0.011	0.004	0.003	0.003	0.003	0.002	0.002

Table 8: Benchmarking against realized beta

This table provides the time-series averages of the cross-sectional correlation coefficient (CORR, Eq. (16)) and the cross-sectional root mean squared error (RMSE, Eq. (17)) of all estimated betas benchmarked against realized beta (Eq. (15)). While the estimated betas are based on historical data prior to month t, realized beta is estimated in-sample using daily returns in month t. Panels A, B, C, and D cover OLS betas (Section 4.2), adjusted 24-month daily regression-based betas (Section 4.3), option-implied betas by Buss and Vilkov (2012) (Section 4.4), and portfolios betas (Section 4.5), respectively. The rows Rank(CORR), Rank(RMSE), and Rank(γ) report the relative ranking of each beta based on CORR, RMSE, and γ, respectively.

Panel A: OLS betas											
	OLS_D^{12}	OLS_D^{6}	OLS_D^{24}	OLS_D^{36}	OLS_D^{60}	OLS_M^{12}	OLS_M^{24}	OLS_M^{36}	OLS_M^{60}		
CORR	0.364	0.345	0.362	0.356	0.341	0.170	0.206	0.218	0.224		
RMSE	1.173	1.206	1.167	1.167	1.174	1.772	1.485	1.405	1.350		
Rank(CORR)	1	4	2	3	5	9	8	7	6		
Rank(RMSE)	3	5	1	2	4	9	8	7	6		
Rank(γ)	4	6	1	2	3	9	8	7	5		
Panel B: Adjusted regression-based betas											
	OLS_D^{24}	BLM	BSW	$BSWA$	VCK	$EWLS_{(84)}$	$EWLS_{(168)}$	DIM	SW	FP	AVG
CORR	0.362	0.362	0.370	0.384	0.369	0.389	0.385	0.279	0.327	0.313	0.376
RMSE	1.167	1.166	1.157	1.148	1.158	1.145	1.148	1.263	1.212	1.237	1.156
Rank(CORR)	7	7	5	3	6	1	2	11	9	10	4
Rank(RMSE)	8	7	5	3	6	1	2	11	9	10	4
Rank(γ)	10	11	8	6	9	7	5	3	4	2	1
Panel C: Option-implied betas											
	OLS_D^{24}	BV_{30}	BV_{91}	BV_{182}	BV_{273}	BV_{365}	BV_{AVG}				
CORR	0.568	0.543	0.584	0.590	0.591	0.588	0.587				
RMSE	0.567	0.569	0.550	0.550	0.551	0.554	0.550				
Rank(CORR)	6	7	5	2	1	3	4				
Rank(RMSE)	6	7	3	1	4	5	2				
Rank(γ)	7	6	4	1	3	5	2				
Panel D: Portfolio betas											
	OLS_D^{24}	$P2$	$P5$	$P10$	$P25$	$P50$	$P75$	$P100$			
CORR	0.362	0.294	0.348	0.359	0.363	0.364	0.364	0.363			
RMSE	1.168	1.187	1.166	1.162	1.160	1.160	1.160	1.160			
Rank(CORR)	5	8	7	6	3	1	2	4			
Rank(RMSE)	7	8	6	5	3	1	2	4			
Rank(γ)	2	8	7	6	5	1	3	4			

Table 9: γ estimates - Mutual Fund Betas

This table reports estimates of γ for a sample of actively managed mutual funds. The sample spans August 1999 to December 2023 and includes 3,098 unique funds. Panel A presents OLS and Vasicek-adjusted (BMVCK) betas estimated over 60- and 36-month rolling windows using monthly returns, as well as over 12- and 6-month windows using daily returns. Panel B reports beta estimates based on exponentially weighted least squares (EWLS) over a 12-month horizon, using half-life parameters of 84 and 168 days, respectively. It also includes Dimson-adjusted betas constructed with four return lags, following the same specification as in the stock-level analysis. The final column reports the simple average across all ten beta estimators, denoted as AVG. $\Delta\gamma$ is measured relative to the benchmark OLS_M^{60}.

Panel A: OLS and Vasicek-adjusted Beta

	OLS_M^{60}	$BMVCK^{60}$	OLS_M^{36}	$BMVCK^{36}$	OLS_D^{12}	$BVCK^{12}$	OLS_D^{6}	$BVCK^{6}$
γ	0.167	0.168	0.163	0.167	0.177	0.178	0.180	0.180
SE[γ]	0.020	0.019	0.021	0.019	0.022	0.021	0.023	0.023
$\Delta\gamma$	.	0.001	-0.004	0.000	0.010	0.011	0.013	0.014
SE[$\Delta\gamma$]	.	0.002	0.002	0.002	0.010	0.010	0.011	0.011

Panel B: EWLS, Dimson, and Simple Average Beta

	$EWLS_{(84)}$	$EWLS_{(168)}$	$DIM_{(12)}$	$DIM_{(6)}$	AVG
γ	0.180	0.179	0.195	0.194	0.191
SE[γ]	0.023	0.022	0.021	0.022	0.021
$\Delta\gamma$	0.014	0.012	0.028	0.027	0.024
SE[$\Delta\gamma$]	0.011	0.011	0.009	0.009	0.007

Table A.1: Unbiasedness of the OLS estimator of γ

Panel A reports rejection frequencies for the hypothesis $H_0 : \phi_1 = 0$ in regression (A.6). For each stock, we run a time-series regression of market returns $R_{m,t}$ on estimated betas $\widetilde{\beta}_{i,t}$, and test the hypothesis that the slope is equal to zero using a t-test. The sample is restricted to stocks with at least 24 monthly observations of $\widetilde{\beta}_{i,t}$, which gives a sample N of 13,015 stocks. The frequencies below report how often, out of these N time-series regressions, H_0 is rejected at the 10%, 5%, and 1% level, for each of the regression-based beta estimators. Panel B reports the slope coefficient ψ_1 and Newey-West standard error from the time-series regression (A.7), for each of the regression-based beta estimators.

Panel A1: OLS betas										
a	OLS_D^{12}	OLS_D^{6}	OLS_D^{24}	OLS_D^{36}	OLS_D^{60}	OLS_M^{12}	OLS_M^{24}	OLS_M^{36}	OLS_M^{60}	
10%	0.113	0.110	0.124	0.097	0.101	0.120	0.139	0.123	0.114	
5%	0.058	0.057	0.070	0.054	0.058	0.063	0.081	0.066	0.061	
1%	0.012	0.012	0.017	0.013	0.016	0.017	0.023	0.016	0.018	
Panel A2: Adjusted regression-based betas										
a	BLM	BSW	$BSWA$	VCK	$EWLS_{(84)}$	$EWLS_{(168)}$	DIM	SW	FP	AVG
10%	0.124	0.129	0.147	0.129	0.150	0.148	0.156	0.135	0.079	0.140
5%	0.070	0.074	0.083	0.074	0.083	0.084	0.088	0.074	0.041	0.077
1%	0.017	0.018	0.022	0.018	0.020	0.022	0.023	0.018	0.009	0.018
Panel B1: OLS betas										
a	OLS_D^{12}	OLS_D^{6}	OLS_D^{24}	OLS_D^{36}	OLS_D^{60}	OLS_M^{12}	OLS_M^{24}	OLS_M^{36}	OLS_M^{60}	
ψ_1	-0.075	-0.008	-0.015	-0.010	-0.006	-0.004	-0.010	-0.007	-0.003	
SE[ψ_1]	0.074	0.012	0.029	0.035	0.037	0.003	0.007	0.010	0.011	
Panel B2: Adjusted regression-based betas										
ψ_1	-0.019	-0.015	-0.021	-0.017	-0.027	-0.021	-0.021	0.021	0.004	-0.024
SE[ψ_1]	0.038	0.040	0.039	0.033	0.037	0.039	0.021	0.020	0.016	0.046