Beta Horizons*

Paul Karehnke† Frans de Roon‡

March 6, 2024

Abstract

We relate beta estimates at different frequencies to each other. Both the unconditional distribution of betas and the distribution conditional on their monthly estimate become more dispersed and closer to a lognormal distribution as the estimation horizon increases. Longer horizon beta estimates that assume either independently and identically distributed returns or simple autocorrelation dynamics provide much better estimates of longer horizon betas than short-horizon betas. Longer horizon betas and mean holding period returns yield positive and reasonable market risk premium estimates, even for characteristics-sorted portfolios that have been shown to pose a challenge for the CAPM.

Keywords: beta, long-horizon, cost of capital

JEL Classification: G10, G12, G31

*We thank seminar participants at Neoma and the 7th Asset Pricing Breakfast as well as Scott Cederburg for helpful comments and discussions.

†ESCP Business School, Department of Finance, 6-8 avenue de la Porte de Champerret, 75017 Paris, France. Email: pkarehnke@escp.eu.

‡TIAS School for Business and Society and Tilburg School of Economics and Management (TiSEM), Tilburg University, The Netherlands. Email: F.A.deRoon@tilburguniversity.edu.

1 Introduction

Beta as a measure of systematic risk is one of the key ingredients in asset pricing, corporate finance, and many other related fields. Although it has its origins in the Capital Asset Pricing Model (CAPM) as a measure of an asset's risk viz-a-viz the market portfolio return, multifactor models such as the Fama and French (1993, 2015) three- and five-factor models employ similar risk measures—regressing asset returns on factor portfolios such as the size and the value factor. Key in understanding the cross-section of expected stock returns, beta measures are equally important in many practical applications such as performance measurement and determining discount rates for company valuation and investment decisions.

Although empirically betas are usually estimated using relatively short horizon (high frequency) returns—say monthly—their application is often for longer horizon (lower frequency) expected returns. For instance, in capital budgeting or valuations, one typically uses annual expected cash flows and discounts these using annualized discount rates. In performance evaluation, alphas are typically required for annualized returns, but based on monthly observations. To the extent betas estimated at different horizons coincide, we can limit ourselves to properly annualizing interest rates, equity premiums and other factor risk premiums. However, if betas do not translate one to one from one horizon to another, this requires additional adjustments when using for instance annual expected returns or discount rates estimated from monthly returns.

Our task in this paper is a simple one and perhaps rather mundane. We analyze the distribution of betas across different return horizons. To this end, we use a bootstrapped sample of 4,600 stocks for 36,000 months. Based on this sample, we calculate simple CAPM-type betas for horizons of one month up to five years. Our findings show that longer horizon betas differ substantially from monthly ones. Whereas the distribution of monthly betas appears to be rather close to a normal distribution (mean one, standard deviation close to one half; at least for the left side), as the horizon increases, the empirical distribution starts to resemble a lognormal distribution more and more. Indeed, the difference in betas occurs especially in the right half of the distribution. For instance, going from monthly to annual betas, the 90%-percentile beta increases from 1.75 to 2.23, more than 25% higher. Assuming an equity premium of five percent, this difference of about 0.5 translates into a difference in cost of capital of 2.5% per year. Such differences, although smaller in the extreme percentiles, also occur for the distribution of longer horizon betas conditional on the monthly betas.

Predicting the longer-horizon betas conditional on the monthly betas is the second goal of this paper. Previous literature has shown the exact formulas for calculating different horizon betas when returns are independently and identically distributed (iid; e.g., Levhari and Levy, 1977, Handa, Kothari, and Wasley, 1989, Bessembinder, Cooper, and Zhang, 2024). Based on Stuart and Ord (1994, Chapter 10), we use a Taylor series approximation of returns around their expected value to obtain a simple expression for longer horizon betas conditional on a shorter horizon estimate. This also allows to come up with expressions for longer horizon betas in case returns are not iid, but can be described by a vector autoregressive (VAR) model or a univariate first-order autoregressive (AR(1)) model, for example. Regressing longer horizon betas from our bootstrapped sample on candidate betas, we find that betas implied by assuming either iid returns or univariate AR(1) processes yield intercepts close to zero and slope coefficients close to one, and similarly so for more complex VAR-models. However, especially the behavior of betas in the left tail of the distribution cannot be captured by the simple iid or AR(1) processes, but require the dynamics implied by a VAR model.

As a third and last objective of our paper, we study the performance of the CAPM using mean holding period returns and betas for different horizons. For our 4,600 individual bootstrapped stocks, cross-sectional regressions provide reasonable estimates of slope coefficients close to the in-sample equity premium and intercepts around 50% larger than the average risk-free rate. For holding period returns of two years and betas estimated with two year returns, the estimates are most reasonable. At the same time, in terms of explaining the cross-sectional variation in average returns, the CAPM performs best with average five-year returns and five-year betas. In this case, the R^2 is 54% versus 16% at the one-year horizon.

We find similar horizon effects for characteristics-sorted portfolios. Most notably, we show that increasing the return and beta horizon reduces the intercepts in cross-sectional regressions and changes the slope coefficients from significantly negative or flat as found in prior studies (see Fama and French, 1992; Jagannathan and Wang, 2007, among others) to significantly positive. The estimates for the market risk premium are thereby reasonable and mostly in the range of 5-8% per year for the ten-year beta horizon for the 25 portfolios formed on size and book-to-market (B/M), size and momentum, and operating profitability and investment.

Our analysis of long horizon betas using a bootstrapped sample is similar in nature to the analysis in Fama and French (2018a,b) on long horizon mean returns and volatilities, Anarkulova, Cederburg, and O'Doherty (2022, 2023) on international stock and bond returns, and Aretz and Arisoy (2023) on long horizon skewness. We study the cross-section of stock betas effectively focusing on the joint distribution of stocks and the market portfolio, rather than returns at the stock or market portfolio level. We also deviate from Fama and French (2018a,b) and focus, as Anarkulova et al. (2022, 2023) and Aretz and Arisoy (2023), on the dynamics of returns as they deviate from being iid.

Our paper adds most directly to the growing literature on beta horizon effects. For example, Gilbert, Hrdlicka, Kalodimos, and Siegel (2014) study daily, monthly, and quarterly betas and conclude that "daily data is bad for beta" because at high frequencies the impact of systematic news is revealed only with a delay. Levi and Welch (2017) study best practices for capital budgeting and emphasize the importance of applying Vasicek (1973) shrinkage to daily beta estimates. Our study focuses on very long horizon betas and the shortest horizon that we consider is monthly. Our approximate formulas for horizon betas that allow for non-iid returns add to the exact horizon formulas derived by Levhari and Levy (1977), Handa et al. (1989), and Bessembinder et al. (2024).[1] A related strand of the literature considers multiple frequencies to test asset pricing models (e.g. Beber, Driessen, Neuberger, and Tuijp, 2020; Brennan and Zhang, 2020; Chernov, Lochstoer, and Lundeby, 2022; Kamara, Korajczyk, Lou, and Sadka, 2016; Kothari, Shanken, and Sloan, 1995), though none of those papers document a positive price of very long-horizon market risk for characteristics-sorted portfolios.

The rest of the paper proceeds as follows. In the next section, we explain our bootstrapped sample and various ways to calculate implied longer horizon betas from short horizon ones. Section 3 analyzes the unconditional and conditional distribution of betas across different return horizons. Section 4 further investigates the cross-sectional behavior of betas and stock returns. Section 5 concludes.

2 Data and methodology

2.1 Data

Our basic data set consists of monthly individual stock returns for the period January 1962 to December 2022 from the Center for Research in Securities Prices (CRSP) of all common shares (share code 10 or 11) listed on the NYSE, Nasdaq, and AMEX. Next to individual stock returns, we use the value weighted market return as a proxy for our market index and the risk-free rate, when excess returns are needed, as provided on Kenneth French's website.[2]

[1] In the conclusion of their study on mutual fund performance evaluation, Bessembinder et al. (2024) remark (see page 27): "While the estimates we obtain appear reasonable, they do rely on simplifying assumptions (including that returns are iid over time) that are violated in the data, and we anticipate that future researchers may well be able to refine these methods." We provide such formulas that allow for non-iid returns.

[2] See http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html.

Denoting the monthly total stock return for stock i as $R_{i,t+1} = R_{i,t\to t+1}$, we construct K-month holding period returns ($K = 1, 2, ..., 60$) as

$$1 + R_{i,t\to t+K} = \prod_{s=1}^{K} (1 + R_{i,t+s}) \, . \tag{1}$$

These K-month holding period returns are the object of our analysis.

2.2 The bootstrap

We follow the stationary bootstrap proposed by Politis and Romano (1994) and described by, for example, Sullivan, Timmermann, and White (1999) and Anarkulova et al. (2022), adjusted to our setting. First, we select all available stocks for which there are at least ten years of monthly returns available. The total available time-series of these stocks is our universe of eligible stock-month observations.

Next, out of this universe of eligible observations, we randomly select 4,600 different stocks (with replacement) to mimic the number of available stocks in the median month in our sample period. For each stock, we then select blocks of returns with an average random block length of six months. In unreported results, we verify that this random block length sufficiently captures the dependencies in returns that drive our results. We repeat this block sampling for each of the 4,600 stocks until we have a total series of 36,000 monthly returns. Our main analysis is performed on this bootstrapped sample.

2.3 Descriptive statistics

In Table 1, we show that the stocks selected for our bootstrap analysis (Panel A) are similar to the universe of eligible stocks (Panel B) in terms of number of available time-series observations, average returns, standard deviations, and betas. Comparing the universe of eligible stocks to all available stocks (Panel C) reveals that stocks with shorter return histories that are excluded from our bootstrap analysis tend to have rather extreme average returns (mostly low), extreme volatilities (mostly high), and extreme betas (mostly low). However, except for the much wider dispersion of betas in the universe of all available stocks (standard deviation of 1.37 vs. 0.51, reflecting fatter tails), the cross-sectional distribution is rather similar to the one in the universe of eligible stocks and our sample of stocks for the bootstrap. For example, the means and medians are very similar (1.10 vs. 1.11 to 1.12 and 1.05 vs. 1.08) and the 10th and 90th percentiles are a bit wider than the 5th and 95th percentiles in the subsamples (0.10 vs. 0.36 and 2.11 vs. 2.00). In sum, Table 1 suggests that the horizon effects we document in Section 3 are likely representative for the majority of the firms in the CRSP file.

2.4 Beta comparisons

Our starting point is beta calculated from returns over different holding period horizons K, using the returns as defined above:

$$\beta\left(K\right)_{im} = \frac{\text{Cov}\left[R_{i,t\rightarrow t+K}, R_{m,t\rightarrow t+K}\right]}{\text{Var}\left[R_{m,t\rightarrow t+K}\right]}. \tag{2}$$

For monthly returns, we will sometimes simply write $\beta\left(1\right)_{im} = \beta$. We refer to the estimated $\beta\left(K\right)_{im}$ in our bootstrapped sample as the empirical betas.

Implied betas for the iid case

When returns are iid, we can use a Taylor series expansion of $R_{i,t+s}$ and $R_{m,t+s}$ around their expected values μ_i and μ_m similar to Stuart and Ord (1994, Chapter 10) to obtain:

$$\text{Cov}\left[R_{i,t\rightarrow t+K}, R_{m,t\rightarrow t+K}\right] = K\left(1+\mu_i\right)^{K-1}\left(1+\mu_m\right)^{K-1}\sigma_{im}, \tag{3}$$

with $\sigma_{im} = \text{Cov}\left[R_{i,t+1}, R_{m,t+1}\right]$, the one-period (monthly) covariance. It follows readily that the implied iid beta equals:

$$\beta\left(K\right)_{im}^{\text{iid}} = \frac{\left(1+\mu_i\right)^{K-1}}{\left(1+\mu_m\right)^{K-1}}\beta_{im}. \tag{4}$$

Note that exact formulas for this beta are provided by, for example, Levhari and Levy (1977), but we find this Taylor series approximation to work very well and it can easily be extended to allow for autocorrelation in returns as we show below. We provide details of our derivations in Appendix A.

When applying (4), we need to estimate the one-period expected returns μ_i and μ_m. As a first estimate, we use the simple average (monthly) returns in our bootstrapped sample. this results in $\beta\left(K\right)_{im}^{\text{iid}}$. As a second estimate, we impose the CAPM, such that in the simulated sample we have

$$\beta\left(K\right)_{im}^{\text{CAPM}} = \frac{\left(1+\mu_i\right)^{K-1}}{\left(1+\mu_m\right)^{K-1}}\beta_{im}, \tag{5}$$

where

$$\mu_i = R_f + \beta_{im}\left(\mu_m - R_f\right). \tag{6}$$

Implied betas with autocorrelations in returns

As a next step, we relax the assumption of iid returns in the approximation and assume that $\boldsymbol{R}_{t+1}^{\top} = \begin{bmatrix} R_{i,t+1} & R_{m,t+1} \end{bmatrix}$ follow a first-order vector autoregressive (VAR) process:[3]

$$\begin{pmatrix} R_{i,t+1} \\ R_{m,t+1} \end{pmatrix} = \begin{pmatrix} \mu_i \\ \mu_m \end{pmatrix} + \boldsymbol{\Gamma}\begin{pmatrix} R_{i,t} - \mu_i \\ R_{m,t} - \mu_m \end{pmatrix} + \begin{pmatrix} u_{i,t+1} \\ u_{m,t+1} \end{pmatrix}, \tag{7}$$

[3] Such a VAR process can be motivated by Boguth, Carlson, Fisher, and Simutin's (2016) analysis with heterogeneous information diffusion and leading and lagging stocks.

where

$$\mathbf{\Gamma} = \begin{bmatrix} \gamma_{ii} & \gamma_{im} \\ \gamma_{mi} & \gamma_{mm} \end{bmatrix}. \tag{8}$$

Defining

$$\mathbf{\Omega}(K) = \left\{ \frac{1}{2} K \times \boldsymbol{I}_K + \sum_{j=1}^{K-1} (K-j)\, \mathbf{\Gamma}^j \right\} \text{Var}\left[\boldsymbol{R}_{t+1}\right], \tag{9}$$

we can write

$$\text{Cov}\left[R_{i,t\to t+K}, R_{m,t\to t+K}\right] = (1+\mu_i)^{K-1} (1+\mu_m)^{K-1} \left\{\Omega(K)_{im} + \Omega(K)_{mi}\right\}, \tag{10}$$

where Ω_{im} is the element in the first row and second column of the 2×2 matrix $\mathbf{\Omega}$. By choosing i and m for the numerator and $i = m$ for the denominator, the implied beta follows readily. We refer to these betas as $\beta(K)_{im}^{\text{VAR}}$.

Suppose $\mathbf{\Gamma}$ is diagonal, so there are no cross-autocorrelations. In that case, we have

$$\mathbf{\Omega}(K) = \left\{ \frac{1}{2} K \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} + \sum_{j=1}^{K-1} (K-j) \begin{bmatrix} \gamma_i^j & 0 \\ 0 & \gamma_m^j \end{bmatrix} \right\} \begin{bmatrix} \sigma_{ii} & \sigma_{im} \\ \sigma_{mi} & \sigma_{mm} \end{bmatrix} \tag{11}$$

and

$$\mathbf{\Omega}(K)_{im} + \mathbf{\Omega}(K)_{mi} = \left\{ K + \sum_{j=1}^{K-1} (K-j) \left(\gamma_i^j + \gamma_m^j\right) \right\} \sigma_{im}. \tag{12}$$

Using this restricted form in (10) gives us the implied beta when stock and market returns each follow an AR(1) process with autocorrelation parameters γ_i and γ_m, respectively. The betas implied by these univariate AR(1) processes will be referred to as $\beta(K)_{im}^{\text{AR}}$.

3 Empirical and implied betas for longer horizons

3.1 The unconditional distribution of betas

Table 2 shows our first main results. The table shows the distribution of empirical betas, as in equation (2), for horizons K varying from one month to 60 months. These distributions are referred to as the unconditional distributions. For monthly returns ($K = 1$), the median beta is 1.09 and the left side of the distribution is well approximated by a normal distribution with mean one and standard deviation 0.4. For instance, for such a normal distribution the 25% percentile would be 0.73 (0.77 in our sample), the 10% percentile would be 0.49 (0.50 in our sample), and the 5% percentile would be 0.34 (0.36 in our sample). For the right-hand side of the distribution, this is different: although the 75% percentile is still rather close (1.27 in the normal distribution versus 1.43 in the sample), the further we go into the tail the worse the normal approximation is. For the 90% percentile the numbers are 1.51 versus 1.75 and for the 95% percentile they are 1.66 versus 1.99. Thus, the empirical distribution of monthly betas appears to be more right-skewed than the normal distribution.

The remaining rows of Panel A of Table 2 show the distribution of empirical betas for longer holding period returns (horizons K). It is clear from Panel A that the distribution of betas changes as the estimation horizon K increases. The median of the distribution changes only slightly from 1.09 for a one month horizon to about 1.18 at the annual horizon and 1.24 at the five-year horizon. For all percentiles below 50%, the betas tend to systematically decrease with the horizon K, whereas for all percentiles above 50% they systematically increase. At a horizon of one year, the 90% percentile shows a beta that is almost 0.5 (27%) higher than the monthly one (2.23 versus 1.75), and at the 10% percentile the difference is 0.05 (10%). With an equity premium of five percent, a difference in beta of 0.5 translates into a sizable difference in cost of capital of 2.5% per year. As the horizon increases beyond one year, the differences with the monthly estimates become more extreme, especially for the right-hand side of the distribution.

As the horizon K increases, the distribution of empirical betas in Panel A also appears to deviate more and more from the normal distribution and starts to resemble a lognormal distribution. This finding, and the fact that below-median betas decrease with the horizon, whereas above-median betas increase with the horizon, are in line with the approximation in equations (4) and (6), at least when high (low) beta stocks are linked to high (low) mean returns. From equations (4) and (6), it follows immediately that the distribution of betas becomes more disperse with the holding period and the high (low) beta/mean return stocks have increasingly higher (lower) betas.

3.2 The conditional distribution of betas

Panel B of Table 2 shows the distribution of empirical betas, conditional on the beta from monthly returns ($K = 1$). To be precise, for each percentile reported in Table 2, we keep the stocks that are within ±0.5% of that percentile and estimate the empirical betas for longer horizons. For instance, for the 75% percentile, we estimate the longer horizon betas for all stocks that are in the 74.5%-75.5% percentile and report the average betas of these stocks. Thus, for each percentile, the different rows show the average beta for a holding period K, conditional on the one-month beta being in a specific percentile q ($\pm 0.5\%$).

We first observe that the median beta (50% percentile) increases with the horizon K. The pattern we observe is qualitatively consistent with a market beta equal to one and a monthly median beta of 1.09, but not quantitatively since the equity premium we need to match it would be six percent per month and thus too high. In the next section, we show that the horizon betas calculated with a VAR(1) are able to yield the pattern we observe for the median beta.

Second, looking at the extreme percentiles ($q = 1\%, 5\%, 95\%, 99\%$), we see that the distribution of betas for longer horizons becomes less dispersed compared to Panel A. This is a form of regression to the mean: conditioning on the monthly beta being at a specific percentile, we see fewer extreme betas for longer horizons. The low percentiles for monthly betas, $q = 1\%, 5\%, 10\%, 25\%$, are somewhat puzzling at first sight, as the betas tend to increase with the horizon for this percentile. This cannot be understood from equations (4) or (6), and as we will show below we need more complicated return dynamics as in the VAR system of equation (7) to generate this type of pattern.

Finally, we observe again that the distribution of betas appears to be close to a lognormal distribution as the horizon increases. For instance, the percentiles of the five year betas in Panel B are well approximated by a lognormal distribution with a (log) mean of 0.2 and a (log) standard deviation of 0.7.

The main interest in this paper is to estimate long horizon betas from short horizon (e.g., monthly) betas—as this is the typical issue faced by analysts and researchers: deriving long term discount rates from short term beta estimates. The betas as presented in Panel B of Table 2 are therefore the primary focus in what follows.

3.3 Explaining long horizon betas

We next proceed to analyze whether the different long horizon betas as discussed in Section 2.4 can explain the observed behavior of betas in Panel B of Table 2. To save space, we only present results that compare betas at the annual frequency conditional on those observed at the monthly frequency. To this end, Panel A of Table 3 first presents three sets of empirical (bootstrapped) betas. The first one shows the monthly betas as in Panel B of Table 2 for the different percentiles. The second row presents the annual betas as in Panel B of Table 2, but now for a bootstrapped sample with block length equal to one. This is a sample that is by construction iid. The third row reports the same empirical annual beta, but now based on our main bootstrapped sample with average block length equal to six months. This row equals the annual conditional betas in Panel B of Table 2 and is our main target for the various approaches to calculate horizon betas.

As Panel A of Table 3 shows, the fact that returns are not exactly iid (the difference between the two bootstrapped samples underlying the second and the third row) creates differences in the behavior of longer horizon betas—in this case annual ones. Overall, the non-iid sample is more right skewed. This is visible in the left side of the distribution ($q < 50\%$) that is less dispersed.

Panel B of Table 3 tries to match the observed conditional distribution of annual betas in Panel A using the annual betas as implied by different return distributions. The first row of Panel B uses equation (4), which assumes that returns are iid. The observed implied betas, $\beta(K)^{\text{iid}}$, closely match the annual betas in our iid sample in Panel A. This confirms that the approximation underlying equation (4) works well. As the implied iid betas are almost identical to the empirical ones in the iid sample in Panel A, they obviously fall short in explaining the empirical annual betas in our non-iid sample in Panel A: the implied iid betas do not capture the left side of the empirical distribution, nor the pattern in the median and 75 percentile betas.

The second row of Panel B aims to capture the return dynamics by assuming independent AR(1) processes for the individual stock returns and the market returns. That is, in the VAR-system in equation (7) the cross-autocorrelations are assumed to be zero: $\gamma_{im} = \gamma_{mi} = 0$. The AR(1) coefficients for the stocks and the market are estimated over the same sample as the monthly betas and mean returns are calculated. At first sight, the betas implied by the AR(1) returns appear to perform worse then the implied iid betas in the first row of Panel B. Ignoring the lowest percentiles ($q \leq 10\%$), comparing the implied AR(1) betas with the implied iid ones shows that the AR(1) betas systematically underestimate the iid ones by about 4 to 10 percent. Dividing the AR(1) beta by the empirical iid one gives 0.90-0.96 for every percentile except for $q = 25\%$ (0.87).

We next analyze whether the full VAR model as in equation (7), that also allows for non-zero cross-correlations, may do a better job in explaining annual (and longer horizon) betas. The third row of Panel B of Table 3 shows the distribution of these betas. These results show that the full VAR model performs best, although it tends to overestimate the betas in the highest percentiles. Overall the absolute difference between the implied VAR betas and the empirical ones in Panel A hardly exceed ten percent. Also, looking across the percentiles, the differences between the implied VAR betas and the empirical ones are particularly small in the left tail. For example, the implied VAR beta at the first percentile is 0.27 versus 0.28 empirically.

Finally, Panel C adds a CAPM restriction to the implied iid and VAR betas. We thus impose that $\mu_i = R_f + \beta_{im} (\mu_m - R_f)$. For the implied iid betas, adding the CAPM restriction mostly seems to deteriorate the predictive ability for annual betas, although the differences are not that large. For the implied VAR betas, the CAPM restriction improves the performance, as overall the differences between the implied and empirical betas become smaller, especially in the extreme percentiles.

We next proceed by analyzing the explanatory ability of the different specifications in a more systematic way. Specifically, given a candidate longer horizon beta, $\beta(K)^{\text{cand}}$, we estimate the regression:

$$\beta(K)_{im} = a + b\beta(K)_{im}^{\text{cand}} + e_i, \text{ for } i = 1, ..., 4600. \tag{13}$$

If we explain the longer horizon betas well, we should find $a = 0$ and $b = 1$. In addition, we should have a high value for the R^2. Table 4 reports the results. As a benchmark, the first three rows first report the estimates for a naive approach in which the candidate beta is simply the one-month beta; that is, $\beta(K)^{\text{cand}} = \beta(1)$. We next report the estimates for the different approaches in the same order as in Table 3.

Looking across models, we see that R^2s decrease with the horizon K. In terms of R^2 the best performing model is the VAR(1) model. Imposing the CAPM restriction in the last two candidate betas in Table 4, the R^2s decrease relatively fast with the horizon K. Considering the estimated intercept and slope coefficients a and b in Table 4 suggests that the AR(1) and VAR(1) specifications—with or without the CAPM restriction—generally work best. However, for horizons up to 24 months, the biases in the slope coefficient and the intercept are less than 0.10 for all models except the naive one. For this model, the slope coefficients are generally the furthest from one compared to other specifications, and the intercepts are furthest from zero. For the other models, it is only for the longest horizons that biases increase beyond 0.10, but only for the iid-CAPM model beyond 0.20. Thus, based on these estimates, we do not prefer either of the three models over the others—except that the naive one-month beta is easily beaten by all other models considered. Adding a CAPM restriction to the iid specification increases the bias in both the intercept and the slope coefficient away from zero and one respectively. This effect does not occur with the VAR(1) specification, where the biases are similar when adding the CAPM restriction. For the iid case, the bias from imposing the CAPM restriction is only visible at the longest horizon though, and not severe. The cost of this higher bias may in practice be acceptable compared to the benefit of imposing the CAPM restriction and not having to estimate mean expected returns (which are often the object of interest).

4 Additional results

Having established the difference between monthly and longer horizon betas and ways to approximate the longer horizon betas, we next provide two additional analyses on the cross-section of betas and stock returns. First, we investigate the cross-sectional relation between betas and mean returns at different horizons. Next, we analyze whether long horizon betas are better able to explain average returns on characteristics-sorted portfolios that have been notoriously difficult to capture by the CAPM at monthly horizons.

4.1 Cross-sectional regressions

As our analysis so far suggests that betas for different holding periods differ in their distribution and are all but perfectly correlated, they may contain different information on average returns in the cross-section of stocks. As mentioned earlier, and a prime motivation for this paper, the CAPM is a one period model but is silent about the length (horizon) of one period. It may therefore very well be that the CAPM fails to explain the relation between average returns and beta at one horizon, whereas it may be more successful at other horizons.

To address this question, we further exploit our bootstrapped sample by running cross-sectional regressions:

$$\overline{R}_{i,t\rightarrow t+K} = \lambda_{K,L,0} + \lambda_{K,L,1}\beta\left(L\right)_{im} + u_i. \tag{14}$$

We thus allow the horizon to be different for the dependent variable, the average holding period return $\overline{R}_{i,t\rightarrow t+K}$, and for the estimation horizon of the independent variable, $\beta\left(L\right)_{im}$. Table 5 reports the results, with the annualized estimated constant term in Panel A, the annualized estimated slopes in Panel B, and the R^2s in Panel C. Again, as we have a very big bootstrapped sample (4,600 individual stocks), we refrain from reporting t-statistics, but in this sample all coefficients are at least eight standard errors away from zero, and except for two of them even more than ten standard errors.

Panel A shows that the intercept decreases with the holding period return horizon K and first decreases and then increases with the beta horizon L. The highest annualized estimate is $\lambda_{1,60,0} = 11.99\%$, the lowest is $\lambda_{60,24,0} = 5.59\%$. These numbers can be contrasted with the average risk-free rate during the sample period, which is 4.39% per year. Thus, the implied zero-beta rate λ_0 exceeds the risk-free rate and all intercepts are at least one percentage point higher than the average risk-free rate.

In the same vein, the slope coefficients in Panel B increase with the holding period K for the mean returns and first increase and then decrease with the beta horizon L. The annualized estimate of $\lambda_{1,60,1}$, 2.68%, is now the lowest, whereas the annualized estimate of $\lambda_{60,24,1}$, 12.45%, is the highest. These numbers can be contrasted with the average market risk premium of 6.60% during our sample period. Thus, the combination of average 24 months returns with 24 months betas yields the combination of $\lambda_0 = 6.94\%$ and $\lambda_1 = 8.30\%$ that are closest to the sample average risk-free rate (4.39%) and market risk premium (6.60%). The fact that the intercept is off by about 2.5% is in part due to the small firm premium: as our estimates do not differentiate between small and big stocks, small stocks are overweighted in the bootstrapped sample. Note that the small firm premium in our sample is about 2.05%, similar in size to the missed intercept value.

For the R^2s in Panel C, the story is slightly different: here the overall patterns is that betas are best able to explain the cross-sectional variation in mean returns when the estimation horizon of the betas and the holding period of the returns are long. For monthly horizons, the R^2 is 4.1% and this increases to almost 54.08% for the five year horizons. The combination $K = 24$ and $L = 24$, our preferred choice for intercept and slope coefficient, shows a rather modest R^2 of 30.25%. On the other hand, for the five year horizons, $K = L = 60$, which yield the highest R^2, the annualized intercept and slope are $\lambda_0 = 8.44\%$ and $\lambda_1 = 8.54\%$, both within reach of the time series averages discussed above, with an additional size premium in the intercept.

We conclude that increasing the horizon, especially for the holding period returns, can in part help improving the performance of the CAPM. For longer horizons, betas are better able to explain the cross-sectional variation in average returns, and the estimated intercept and slope in the cross-sectional regressions are usually more reasonable. Nonetheless, the estimated intercepts are usually still sizable, in particular when compared to the average risk-free rate. This may in part be due to a small firm effect—still leaving a puzzle to the CAPM.

4.2 Characteristics-sorted portfolios

Next, we analyze whether long horizon betas are better able to explain returns on characteristics-sorted portfolios that have been shown to pose a challenge for the CAPM. Since the time-series for those portfolios is sufficiently long to calculate horizon betas, we can directly estimate

$$R_{i,t\rightarrow t+K} = \lambda_{K,L,0,t} + \lambda_{K,L,1,t}\beta(L)_{im} + \upsilon_{i,t\rightarrow t+K} \tag{15}$$

using monthly returns and overlapping data for horizons beyond one month.

We start with value-weighted returns on the familiar 25 size and B/M portfolios retrieved from Kenneth French's website. Table 6 shows the average intercepts and slopes that are annualized (geometric, not arithmetic) and the average adjusted R^2 of the cross-sectional regressions. For beta horizons up to one year ($L = 1$ to 12), we replicate the well-known failure of the CAPM to explain the average returns on the size-B/M portfolios. For example, the average intercept is 15.66% and the average slope is -1.27% in Table 6 for the annual beta and annual return horizon ($L = 12$ and $K = 12$). After subtracting the average annualized risk-free return in our sample period, 4.47%, from the intercept, these values appear very close to those obtained by Jagannathan and Wang (2007) of 11.31% for the intercept and -0.56% for the slope coefficient using non-overlapping annual data in the period from 1954 to 2003. The intercepts we obtain on beta horizons up to one-year are thereby all positive, large, and highly statistically significant, while the slopes are all negative and insignificant. The t-statistics accounting for autocorrelation in parenthesis and accounting for autocorrelation and the errors-in-variables problem in brackets thereby yield the same inference and have a similar magnitude.

Increasing the beta horizon from annual ($L = 12$) to five years ($L = 60$), reduces the intercepts and switches the sign of the slopes from negative to positive. The slopes at the five-year beta horizon are, however, still insignificant and of relatively low magnitude (below 1%). To see whether this pattern accentuates for longer horizons, we also estimate equation (15) for betas calculated using ten-year overlapping data. Those betas yield large and significant slope estimates ranging from 6% to 9%, and the intercepts are around 7%, which is higher but still close to the average risk-free return in our sample period.

To analyze the horizon effects of the size-B/M portfolio betas in greater detail, we present in Figure 1 four plots of horizon betas versus average annual excess return for beta horizons ranging from one month to ten years.[4] These figures clearly show that as the horizon increases, the average excess returns are better explained by their beta, and the slope coefficient in the graph changes from negative to a reasonable positive 7%. As the beta horizon increases beyond one year in Panel (b), the betas of the growth portfolios (e.g., small growth portfolio) decrease from around 1.5 at the annual horizon to around 0.5 at the ten-year horizon in Panel (d). There is a similar tendency of value portfolio betas to increase from around one at the one-month horizon to above one at the ten-year horizon. Such a behavior is consistent with our equation (4) in Section 2.4 in which the betas grow with the horizon when the security's return is larger than the one of the market and decreases with the horizon when it is smaller.

The results in Table 6 and Figure 1 suggest that value portfolios indeed have higher long-horizon betas and growth portfolios have lower long-horizon betas and that the average returns on those portfolios line up with their betas in a way that is more consistent with the CAPM at long horizons. In the remainder of this section, we examine whether this pattern generalizes to the 25 size and momentum portfolios and the 25 operating profitability and investment portfolios from Ken French's website.

Table 7 shows that this is indeed the case. In fact, for those portfolios the slope estimates switch from being significantly negative for beta horizons up to one-year to significantly positive for five-year and ten-year beta horizons. The pattern is most pronounced for the size-momentum portfolios for which the intercepts are even insignificant, but close to the average risk-free rate, at the five year horizon. The scatter plots of horizon betas versus average annual excess returns in Figure 2 document that winner and loser portfolios in fact have betas that are larger than one and of similar magnitude at the one-month and one-year horizon. At the five- and ten-year horizon, however, loser portfolios' betas tend to be far below one, whereas winner portfolios' betas are far above one. Similarly, in Fig-

[4]The annual risk-free rate is from Liu and Wu (2021) available at `https://sites.google.com/view/jingcynthiawu/yield-data?authuser=0`.

ure 3, the beta of the portfolio of weak operating profitability and aggressive investment decreases from above one to below one as the beta horizon increases from one-month to ten-year. The beta of the portfolio of robust operating profitability and conservative investment instead increases from around one at the one-month horizon to around two at the ten-year horizon. In sum, the evidence from this section suggests that average returns on characteristics-sorted portfolios line up with long-horizon betas in a way that is consistent with the CAPM.

4.3 Sub-samples and asset betas

To see whether our results are driven by specific sub-periods or possibly time-variation in betas, we split our sample in two: until June 1992 and starting in July 1992. Table 8 reports the conditional distributions of $\beta(K)$, conditional on $\beta(1)$, as in Panel B of Table 2. To save space, we only report results for $K = 1, 3, 12, 60$. To facilitate comparison, we report the results for the entire sample in Table 2 in Panel A of Table 8. Panels B and C are for the two respective sub-samples.

Focusing first on $\beta(1)$, we see that the distribution of betas is more disperse in the second sub-sample than in the first sub-sample. For the early sample, the 5% and 95% percentiles for $\beta(1)$ are 0.41 and 1.85, respectively, whereas these percentiles are 0.28 and 2.19 in the recent sample. Indeed, the 5% and 95% percentiles of the recent sample are close to the 1% and 99% percentiles of the early sample. The wider dispersion in the recent sample also translates in the 1% percentile $\beta(1)$ being close to zero: 0.03.

As the return horizon K increases, the difference in dispersion in the two sub-samples further increases. Moreover, the behavior of betas in the left and right side of the distributions behave a bit differently in both sub-periods. For the right half, the difference between $\beta(1)$ and the longer horizon betas are usually bigger for the second versus the first part of the sample, especially when we move further into the tail. For the left side on the other hand, the differences are bigger (and more meaningful) in the first part of the sample. For instance, whereas for the recent sample the differences between $\beta(1)$ and $\beta(12)$ are rather small (around 0.08 in absolute value) for the 5% and 10% percentiles, for the early sample these differences are 0.15 (0.41 versus 0.56) and 0.19 (0.54 versus 0.73), respectively. Thus, even though in the iid case we expect the differences in $\beta(K)$ across horizons to be small, this is not what we find for the early sample, implying that more involved dynamics than a simple iid model are needed.

In sum, even though the behavior of the conditional betas across horizons is similar in the total sample and in the two sub-samples, we also observe notable differences in the different sub-samples, in both halves of the distribution.

In the last piece of analysis, we analyze whether the horizon effects we document are present for asset betas that are important for cost of capital calculations. To do this, we calculate unlevered equity returns following Doshi, Jacobs, Kumar, and Rabinovitch (2019) and compute the corresponding conditional beta. This analysis starts in February 1975 due to the more limited availability of leverage data necessary to calculate those unlevered returns. We therefore also calculate the corresponding conditional percentiles for the standard beta using only observations for which unlevered returns are available. Moreover, we compute the conditional percentiles for beta calculated with respect to a market return that is the asset weighted average of unlevered returns. Doshi et al. (2019) use this variant of an asset beta in their asset pricing tests. Table 9 reports the results. While there is no discernible horizon effect in asset betas for the left-hand side of the cross-sectional distribution, the horizon effects on the right-hand side are large and economically significant. In particular, for those percentiles the asset beta in Panel B and Doshi et al.'s (2019) asset beta variant in Panel C both increase significantly with the horizon. For example, in Panel B, asset beta increases by about 0.2 from the one-month to the annual horizon at the 90% percentile (from 1.23 to 1.42). These results suggest that horizon variations in asset betas translate into meaningful cost of capital variations.

5 Conclusions

We show that betas estimated at different holding period horizons (return frequencies) deliver different distributions of betas, which has important implications for estimated costs of capital and performance measures, for instance. Using a block bootstrapped sample of 4,600 stocks and 36,000 months from monthly CRSP data for the period 1962 to 2022, the unconditional distribution of betas becomes more dispersed with the horizon and closer to a lognormal distribution. But also for short horizons; that is, monthly, a normal distribution cannot describe both sides of the distribution well.

Conditioning on the one-month beta, the distribution of betas likewise becomes lognormal at longer horizons and more dispersed. This behavior of the distribution is consistent with a simple Taylor series approximation of longer horizon betas for iid sampled returns, but not for non-iid samples (i.e., with block lengths larger than one). Using a Taylor-series approximation of longer horizon betas assuming either iid returns or returns generated by univariate AR(1) processes give a good first approximation of the longer horizon distribution, but a more involved VAR-system generally performs best.

We use the longer horizon betas to study the performance of the CAPM at different return horizons and find that in cross-sectional regressions longer horizon betas can better explain the cross-sectional variation of longer period holding returns, with reasonable parameter estimates viz-a-viz the sample averages of the market risk premium and the risk-free rate. For characteristics-sorted portfolios, the estimated market risk premium is positive, realistic, and statistically significant for long-horizon betas, while it is negative for short-horizon betas.

A Appendix: derivation of the implied betas

Consider the beta for longer horizons, say K months. Using a Taylor series expansion of $R_{i,t+k}$ and $R_{m,t+k}$ for $k = 1, \ldots, K$ around μ_i and μ_m as in Stuart and Ord (1994, Section 10.5) and assuming that returns are iid, we get

$$\text{Cov}\left[\prod_{k=1}^{K}(1+R_{i,t+k}), \prod_{k=1}^{K}(1+R_{m,t+k})\right] = K(1+\mu_i)^{K-1}(1+\mu_m)^{K-1}\sigma_{im}, \tag{16}$$

where—similar to μ—σ_{im} is the covariance of the one-month returns R_i and R_m. Likewise,

$$\text{Var}\left[\prod_{k=1}^{K}(1+R_{m,t+k})\right] = K(1+\mu_m)^{2(K-1)}\sigma_{mm}. \tag{17}$$

Together, we have for the K-month CAPM $\beta(K)$:

$$\beta(K)_{im} = \frac{K(1+\mu_i)^{K-1}(1+\mu_m)^{K-1}\sigma_{im}}{K(1+\mu_m)^{2(K-1)}\sigma_{mm}} = \frac{(1+\mu_i)^{K-1}}{(1+\mu_m)^{K-1}}\beta_{im} \tag{18}$$

with $\beta(1) = \beta$.

A.1 Autocorrelation in returns

Suppose now that stock and market returns can be described by AR(1) processes:

$$R_{i,t+1} = \mu_i + \gamma_i(R_{i,t} - \mu_i) + u_{i,t+1},$$

$$R_{m,t+1} = \mu_m + \gamma_m(R_{m,t} - \mu_m) + u_{m,t+1}.$$

This means that we have for auto-covariances:

$$\text{Cov}\left[R_{i,t+1}, R_{m,t+k}\right] = \text{Cov}\left[R_{i,t+1}, \rho_m^{k-1} R_{m,t} + \sum_{s=1}^{k-1} \rho_m^s u_{m,t+s}\right] = \rho_m^{k-1}\sigma_{im},$$
$$\text{Cov}\left[R_{m,t+1}, R_{i,t+k}\right] = \text{Cov}\left[R_{m,t+1}, \rho_i^{k-1} R_{i,t} + \sum_{s=1}^{k-1} \rho_i^s u_{i,t+s}\right] = \rho_i^{k-1}\sigma_{im}.$$

Armed with these expressions, and using Section 10.5 from Stuart and Ord (1994) again, we get:

$$\text{Cov}\left[R_{i,t\to t+K}, R_{m,t\to t+K}\right] = \left(1+\mu_i\right)^{K-1}\left(1+\mu_m\right)^{K-1}\left\{K + \sum_{j=1}^{K-1}\left(K-j\right)\left(\rho_i^j + \rho_m^j\right)\right\}\sigma_{im}.$$

This can be written as:

$$\text{Cov}\left[R_{i,t\to t+K}, R_{m,t\to t+k}\right] =$$
$$\left(1+\mu_i\right)^{K-1}\left(1+\mu_m\right)^{K-1}\sigma_{im}\left\{-K + \frac{1}{1-\gamma_i}\left(K - \frac{\gamma_i - \gamma_i^{K+1}}{1-\gamma_i}\right) + \frac{1}{1-\gamma_m}\left(K - \frac{\gamma_m - \gamma_m^{K+1}}{1-\gamma_m}\right)\right\}.$$

A.2 Cross-correlations

If we also allow for cross-correlations, we can use the vector-autoregressive (VAR) model:

$$\boldsymbol{R}_{t+1} := \begin{pmatrix} R_{i,t+1} \\ R_{m,t+1} \end{pmatrix} = \begin{pmatrix} \mu_i \\ \mu_m \end{pmatrix} + \boldsymbol{\Gamma}\begin{pmatrix} R_{i,t} - \mu_i \\ R_{m,t} - \mu_m \end{pmatrix} + \begin{pmatrix} u_{i,t+1} \\ u_{m,t+1} \end{pmatrix}, \qquad (19)$$

where

$$\boldsymbol{\Gamma} = \begin{pmatrix} \gamma_{ii} & \gamma_{im} \\ \gamma_{mi} & \gamma_{mm} \end{pmatrix}. \qquad (20)$$

It seems natural to impose $\gamma_{mi} = 0$ (i.e., individual stock returns will not predict market returns beyond the market itself). The auto-covariance matrices are then given by:

$$\text{Cov}\left[\boldsymbol{R}_{t+1}, \boldsymbol{R}_{t+k}\right] = \Gamma^{k-1}\text{Var}\left[\boldsymbol{R}_{t+1}\right]. \qquad (21)$$

Using

$$\boldsymbol{\Omega}\left(K\right) = \left\{\frac{1}{2}K \times \boldsymbol{I}_K + \sum_{j=1}^{K-1}\left(K-j\right)\boldsymbol{\Gamma}^j\right\} \text{Var}\left[\boldsymbol{R}_{t+1}\right], \tag{22}$$

we can use

$$\text{Cov}\left[R_{i,t\to t+K}, R_{m,t\to t+K}\right] = \left(1+\mu_i\right)^{K-1}\left(1+\mu_m\right)^{K-1}\left\{\Omega_{im} + \Omega_{mi}\right\}, \tag{23}$$

where Ω_{im} is the im-element of the 2×2 autocorrelation matrix $\boldsymbol{\Omega}$, referring to stock i and the market m, or:

$$\boldsymbol{\Omega} = \begin{bmatrix} \Omega_{ii} & \Omega_{im} \\ \Omega_{mi} & \Omega_{mm} \end{bmatrix}. \tag{24}$$

Similarly,

$$\text{Var}\left[R_{m,t+K}\right] = \text{Cov}\left[\prod_{k=1}^{K}\left(1+R_{m,t+k}\right), \prod_{k=1}^{K}\left(1+R_{m,t+k}\right)\right], \tag{25}$$

$$= \left(1+\mu_m\right)^{2(K-1)} 2\Omega_{mm}. \tag{26}$$

Suppose $\boldsymbol{\Gamma}$ is diagonal, so there are no cross-correlations. In that case, we have

$$\boldsymbol{\Omega} = \left\{\frac{1}{2}K\begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} + \sum_{k=1}^{K-1}\left(K-k\right)\begin{bmatrix} \gamma_i^k & 0 \\ 0 & \gamma_m^k \end{bmatrix}\right\}\begin{bmatrix} \sigma_{ii} & \sigma_{im} \\ \sigma_{mi} & \sigma_{mm} \end{bmatrix} \tag{27}$$

and

$$\Omega_{im} + \Omega_{mi} = \left\{K + \sum_{k=1}^{K-1}\left(K-k\right)\left(\gamma_i^k + \gamma_m^k\right)\right\}\sigma_{im}. \tag{28}$$

References

A. Anarkulova, S. Cederburg, and M. S. O'Doherty. Stocks for the long run? evidence from a broad sample of developed markets. *Journal of Financial Economics*, 143(1):409–433, Jan. 2022.

A. Anarkulova, S. Cederburg, and M. S. O'Doherty. Long-horizon losses in stocks, bonds, and bills. Working Paper SSRN, 2023. URL https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3964908.

K. Aretz and Y. E. Arisoy. The pricing of skewness over different return horizons. *Journal of Banking & Finance*, 148:106713, Mar. 2023.

A. Beber, J. Driessen, A. Neuberger, and P. Tuijp. Pricing liquidity risk with heterogeneous investment horizons. *Journal of Financial and Quantitative Analysis*, 56(2):373–408, Apr. 2020.

H. Bessembinder, M. J. Cooper, and F. Zhang. The (large) effect of return horizon on fund alpha. *Critical Finance Review*, forthcoming, 2024.

O. Boguth, M. Carlson, A. Fisher, and M. Simutin. Horizon effects in average returns: The role of slow information diffusion. *Review of Financial Studies*, 29(8):2241–2281, Apr. 2016.

M. J. Brennan and Y. Zhang. Capital asset pricing with a stochastic horizon. *Journal of Financial and Quantitative Analysis*, 55(3):783–827, Oct. 2020.

M. Chernov, L. A. Lochstoer, and S. R. H. Lundeby. Conditional dynamics and the multihorizon risk-return trade-off. *The Review of Financial Studies*, 35(3):1310–1347, Mar. 2022.

H. Doshi, K. Jacobs, P. Kumar, and R. Rabinovitch. Leverage and the cross-section of equity returns. *The Journal of Finance*, 74(3):1431–1471, 2019.

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

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

E. F. Fama and K. R. French. A five-factor asset pricing model. *Journal of Financial Economics*, 116(1):1–22, 2015.

E. F. Fama and K. R. French. Long-horizon returns. *The Review of Asset Pricing Studies*, 8 (2):232–252, 2018a.

E. F. Fama and K. R. French. Volatility lessons. *Financial Analysts Journal*, 74(3):42–53, jul 2018b.

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

T. Gilbert, C. Hrdlicka, J. Kalodimos, and S. Siegel. Daily data is bad for beta: Opacity and frequency-dependent betas. *Review of Asset Pricing Studies*, 4(1):78–117, feb 2014.

P. Handa, S. Kothari, and C. Wasley. The relation between the return interval and betas. *Journal of Financial Economics*, 23(1):79–100, jun 1989.

R. Jagannathan and Y. Wang. Lazy investors, discretionary consumption, and the cross-section of stock returns. *The Journal of Finance*, 62(4):1623–1661, 2007.

A. Kamara, R. A. Korajczyk, X. Lou, and R. Sadka. Horizon pricing. *Journal of Financial and Quantitative Analysis*, 51(6):1769–1793, Dec. 2016. ISSN 1756-6916.

S. P. Kothari, J. Shanken, and R. G. Sloan. Another look at the cross-section of expected stock returns. *The Journal of Finance*, 50(1):185–224, Mar. 1995.

T. A. Kroencke and J. Thimme. A skeptical appraisal of robust asset pricing tests. Working Paper SSRN, 2021. URL `https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3875657`.

D. Levhari and H. Levy. The capital asset pricing model and the investment horizon. *The Review of Economics and Statistics*, 59(1):92, feb 1977.

Y. Levi and I. Welch. Best practice for cost-of-capital estimates. *Journal of Financial and Quantitative Analysis*, 52(2):427–463, Apr. 2017.

Y. Liu and J. C. Wu. Reconstructing the yield curve. *Journal of Financial Economics*, 142 (3):1395–1425, Dec. 2021.

W. K. Newey and K. D. West. A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica*, 55(3):703–708, 1987.

D. N. Politis and J. P. Romano. The stationary bootstrap. *Journal of the American Statistical Association*, 89(428):1303–1313, 1994.

A. Stuart and J. K. Ord. *Kendall's Advanced Theory of Statistics Volume 1 Distribution Theory*, volume 1. Arnold, 6 edition, 1994.

R. Sullivan, A. Timmermann, and H. White. Data-snooping, technical trading rule performance, and the bootstrap. *The Journal of Finance*, 54(5):1647–1691, oct 1999.

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

Table 1:
Descriptive statistics for stocks in bootstrap sample versus all stocks

In this table, we compare the descriptive statistics of the stocks randomly selected for the bootstrap analysis (Panel A) to the universe of eligible stocks (i.e., stocks with at least 120 monthly observations; Panel B) and all available stocks (Panel C). All panels report the average (Mean), standard deviation (Std), and percentiles of the following variables calculated at the stock-level: number of time-series observations T, average return $\bar{R}$, standard deviation σ, and beta calculated with monthly returns β. Panel B and C also report these statistics for the number of stocks with a valid return in each month N. The total number of stocks is 4,600 in Panel A, 9,373 in Panel B, and 25,941 in Panel C. The sample period is from January 1962 to December 2022.

Horizon			**Percentile**								
	Mean	**Std**	**1%**	**5%**	**10%**	**25%**	**50%**	**75%**	**90%**	**95%**	**99%**
Panel A: Stocks used for bootstrap											
T	262.5	138.0	121.0	126.0	133.0	161.0	219.0	316.0	453.0	596.5	732.0
$\bar{R}$	1.35%	0.96%	-1.21%	-0.23%	0.26%	0.90%	1.36%	1.86%	2.37%	2.79%	3.71%
σ	15.74%	7.95%	5.44%	6.93%	7.99%	10.36%	14.22%	19.23%	24.80%	29.26%	41.10%
β	1.12	0.51	0.11	0.36	0.50	0.77	1.08	1.42	1.75	2.00	2.59
Panel B: Stocks with at least 120 monthly observations											
T	262.3	136.9	121.0	126.0	133.0	160.0	219.0	318.0	453.0	583.9	732.0
$\bar{R}$	1.35%	0.97%	-1.37%	-0.27%	0.28%	0.90%	1.35%	1.85%	2.38%	2.79%	3.86%
σ	15.62%	7.77%	5.42%	6.89%	8.01%	10.34%	14.15%	19.10%	24.68%	28.98%	40.10%
β	1.11	0.51	0.11	0.36	0.49	0.76	1.08	1.41	1.76	2.00	2.55
N	3359.0	1028.5	1172.8	1441.5	1688.7	2468.0	3634.5	4210.0	4460.9	4631.9	4770.9
Panel C: All stocks											
T	127.3	133.1	4.0	12.0	17.0	35.0	79.0	172.0	306.0	407.4	633.0
$\bar{R}$	0.25%	4.37%	-15.33%	-7.41%	-4.18%	-0.41%	1.12%	1.98%	3.16%	4.30%	9.10%
σ	18.55%	12.03%	0.58%	6.17%	7.82%	11.18%	16.39%	23.02%	31.05%	37.76%	59.32%
β	1.10	1.37	-1.63	-0.16	0.10	0.58	1.05	1.53	2.11	2.66	4.62
N	4510.4	1500.6	1795.8	2039.0	2093.0	3614.5	4613.0	5707.5	6524.1	6900.4	7379.5

Table 2:
Unconditional and conditional beta distributions

The table shows the percentiles of the cross-sectional distribution of beta as a function of the return horizon. All data is obtained from randomly selecting 4,600 stocks and bootstrapping for each stock a sample of 36,000 monthly returns. Stocks and returns are drawn from common stocks listed on the NYSE, Nasdaq, and AMEX in the CRSP file in the period from January 1962 to December 2022 that have at least ten years of monthly data available. The bootstrap uses the Politis and Romano (1994) approach with random block length equal to six months. Panel A reports the percentiles of empirical betas on the bootstrapped data that are calculated with non-overlapping data for horizons beyond one month. Panel B reports averages of empirical betas in bins taken in the $\pm 0.5\%$ neighborhood of the percentiles of 1m betas. Returns are constructed from monthly returns as $1 + R_{i,t\to t+K} = \prod_{s=1}^{K}(1 + R_{i,t+s})$, for $K = 1, 3, 6, 12, 24, 36, 60$.

Horizon	**Percentile**								
	1%	**5%**	**10%**	**25%**	**50%**	**75%**	**90%**	**95%**	**99%**
Panel A: Unconditional beta distribution									
1m	0.10	0.36	0.50	0.77	1.09	1.43	1.75	1.99	2.60
3m	0.13	0.39	0.54	0.84	1.18	1.57	1.99	2.27	3.08
6m	0.04	0.35	0.50	0.81	1.18	1.64	2.12	2.51	3.37
12m	-0.05	0.27	0.45	0.78	1.18	1.69	2.23	2.72	4.16
2y	-0.20	0.22	0.41	0.74	1.18	1.74	2.51	3.21	5.64
3y	-0.26	0.19	0.37	0.71	1.20	1.86	2.83	3.77	7.45
5y	-0.33	0.14	0.30	0.66	1.24	2.12	3.64	5.52	12.75
Panel B: Distribution of betas conditional on one-month percentile									
1m	0.10	0.36	0.50	0.77	1.09	1.43	1.75	1.99	2.61
3m	0.24	0.44	0.58	0.85	1.17	1.60	1.83	2.21	2.91
6m	0.27	0.46	0.61	0.83	1.17	1.61	1.91	2.26	2.98
12m	0.28	0.50	0.60	0.82	1.20	1.59	1.94	2.28	3.18
2y	0.28	0.52	0.59	0.83	1.23	1.63	2.11	2.49	3.57
3y	0.27	0.54	0.62	0.86	1.29	1.80	2.24	2.68	4.50
5y	0.26	0.56	0.65	0.89	1.49	1.90	2.96	3.29	6.35

Table 3:
Can long horizon betas explain observed betas?

The table analyzes the accuracy of different beta estimates at the annual horizon. Panel A contains empirical betas on the one and twelve month horizon calculated with bootstrapped data using an average block length of one (row 12m iid) and six months (row 1m and 12m). Panel B and C reports implied annual betas that are calculated using the formulas provided in Section 2.4 and with inputs estimated on the bootstrapped sample. Implied iid corresponds to $\beta(12)^{\text{iid}}_{im}$ (see equation (4)), implied VAR(1) is $\beta(12)^{\text{VAR}}_{im}$ (see equation (10)), and implied AR(1) is $\beta(12)^{\text{AR}}_{im}$. CAPM implied iid and VAR betas are calculated as Implied iid and Implied VAR(1) and additionally impose the CAPM to get the expected stock return as in equation (6). The whole table reports averages of these betas in bins taken in the $\pm 0.5\%$ neighborhood of the percentiles of 1m betas (as in Panel B of Table 2).

Beta	**Percentile**								
	1%	**5%**	**10%**	**25%**	**50%**	**75%**	**90%**	**95%**	**99%**
Panel A: Empirical betas									
1m	0.10	0.36	0.50	0.77	1.09	1.43	1.75	1.99	2.61
12m iid	0.09	0.37	0.50	0.76	1.13	1.48	1.96	2.18	3.11
12m	0.28	0.50	0.60	0.82	1.20	1.59	1.94	2.28	3.18
Panel B: Implied betas									
Implied iid	0.10	0.37	0.51	0.76	1.14	1.46	1.96	2.20	3.03
Implied AR(1)	0.09	0.34	0.48	0.72	1.10	1.44	1.87	2.15	2.94
Implied VAR(1)	0.27	0.49	0.64	0.91	1.27	1.68	2.14	2.47	3.27
Panel C: Implied betas with CAPM restriction									
CAPM implied iid	0.09	0.34	0.48	0.75	1.10	1.47	1.86	2.15	2.87
CAPM implied VAR(1)	0.25	0.45	0.60	0.90	1.22	1.70	2.01	2.40	3.10

Table 4:
Explaining long horizon betas
The table reports the estimates of the linear regression $\beta(K)_{im} = a + b\beta(K)_{im}^{\text{cand}} + e_i$ of actual horizon betas on candidate betas and an intercept. Candidate betas are those used in the previous table and the dependent and independent variables are winsorized at the 1% level.

Candidate beta		**Horizon**					
	K	**3**	**6**	**12**	**24**	**36**	**60**
1m	a	0.06	0.04	-0.02	-0.14	-0.26	-0.61
	b	1.06	1.10	1.18	1.35	1.56	2.15
	R^2	83.13%	69.88%	58.33%	47.25%	39.43%	27.40%
Implied iid	a	0.07	0.07	0.06	0.04	0.06	0.14
	b	1.04	1.05	1.04	1.04	1.01	0.96
	R^2	83.21%	70.53%	61.35%	56.93%	55.49%	50.89%
Implied AR(1)	a	0.08	0.08	0.06	0.04	0.04	0.11
	b	1.06	1.08	1.08	1.09	1.07	1.01
	R^2	84.95%	73.04%	64.39%	60.34%	58.99%	53.94%
Implied VAR(1)	a	-0.02	-0.04	-0.05	-0.06	-0.03	0.09
	b	1.01	1.01	1.00	0.98	0.94	0.86
	R^2	95.74%	86.10%	75.81%	68.55%	64.79%	56.18%
CAPM implied iid	a	0.08	0.08	0.08	0.09	0.12	0.20
	b	1.04	1.05	1.06	1.09	1.12	1.22
	R^2	83.16%	70.01%	58.54%	47.52%	39.66%	27.52%
CAPM implied VAR(1)	a	-0.01	-0.03	-0.03	-0.03	0.00	0.07
	b	1.01	1.02	1.02	1.04	1.06	1.16
	R^2	95.67%	85.44%	72.83%	58.66%	48.36%	32.49%

Table 5:
Do horizon betas explain horizon returns?

The table reports the estimates of the cross-sectional regression $\overline{R}_{i,t\to t+K} = \lambda_{K,L,0} + \lambda_{K,L,1}\beta\,(L)_{im} + u_i$ of average horizon returns on horizon betas. The return horizon varies through columns and the beta horizon through rows. The dependent and independent variables in all regressions are winsorized at the 1% level. Panel A and B report the annualized intercepts and slopes, and Panel C reports the adjusted R^2.

Beta \ mean	**K=1**	**3**	**6**	**12**	**24**	**36**	**K=60**	**Average**
			Panel A: Intercept, $\lambda_{L,K,0}$					
L=1	11.86%	11.13%	10.82%	10.79%	10.71%	10.67%	10.59%	10.94%
3	11.54%	10.15%	9.73%	9.67%	9.56%	9.32%	9.16%	9.88%
6	11.42%	10.03%	9.37%	9.19%	9.04%	8.82%	8.52%	9.48%
12	10.89%	9.43%	8.62%	8.16%	7.85%	7.50%	6.88%	8.47%
24	10.41%	8.96%	8.09%	7.52%	6.94%	6.52%	5.59%	7.72%
36	10.53%	9.17%	8.37%	7.82%	7.33%	6.65%	5.88%	7.97%
L=60	11.99%	10.88%	10.25%	9.85%	9.61%	9.42%	8.44%	10.06%
Average	11.23%	9.97%	9.32%	9.00%	8.72%	8.41%	7.87%	9.22%
			Panel B: Slope, $\lambda_{L,K,1}$					
L=1	4.45%	4.83%	5.15%	5.57%	6.73%	8.07%	10.86%	6.52%
3	4.24%	5.09%	5.48%	5.91%	7.02%	8.44%	11.07%	6.75%
6	4.22%	5.06%	5.62%	6.15%	7.28%	8.65%	11.31%	6.90%
12	4.52%	5.40%	6.08%	6.82%	8.06%	9.49%	12.22%	7.51%
24	4.60%	5.44%	6.13%	6.92%	8.30%	9.70%	12.45%	7.65%
36	4.19%	4.91%	5.51%	6.22%	7.48%	8.99%	11.58%	6.98%
L=60	2.68%	3.12%	3.50%	4.00%	4.89%	5.93%	8.54%	4.67%
Average	4.13%	4.84%	5.35%	5.94%	7.11%	8.47%	11.15%	6.71%
			Panel C: R^2					
L=1	4.10%	4.41%	4.54%	4.59%	5.02%	5.36%	5.68%	4.81%
3	5.02%	6.61%	6.94%	6.97%	7.41%	7.96%	8.02%	6.99%
6	6.44%	8.47%	9.46%	9.78%	10.35%	10.89%	10.97%	9.48%
12	10.06%	13.13%	15.03%	16.39%	17.41%	18.14%	18.09%	15.46%
24	17.03%	21.75%	24.94%	27.59%	30.25%	31.13%	30.92%	26.23%
36	22.49%	28.16%	31.97%	35.33%	38.51%	41.75%	40.91%	34.16%
L=60	25.54%	31.54%	35.70%	39.92%	44.05%	46.86%	54.08%	39.67%
Average	12.95%	16.30%	18.37%	20.08%	21.86%	23.15%	24.10%	19.54%

Table 6:
Do horizon betas explain horizon returns for characteristics-sorted portfolios?

The table reports the Fama and MacBeth (1973) estimates of the regression $R_{i,t\to t+K} = \lambda_{K,L,0,t} + \lambda_{K,L,1,t}\beta(L)_{im} + v_{i,t\to t+K}$ of horizon returns on horizon betas for the 25 size and B/M sorted portfolios from Ken French's website. The return horizon varies through columns and the beta horizon varies through rows. The reported average intercepts and slopes are annualized. The calculations use overlapping returns for horizons beyond one month. In parenthesis, we report the t-statistics calculated with Newey and West (1987) standard errors with lag length equal to K. In brackets, we report t-statistics that account for the errors-in-variable problem by using standard errors calculated following Kroencke and Thimme (2021) using the generalized method of moments on the common sample with a lag length equal to the larger of K and L for the Bartlett kernel. The returns are value-weighted. The sample period is from January 1962 to December 2022.

	Intercept, $\lambda_{K,L,0,t}$				**Slope, $\lambda_{K,L,1,t}$**				**Average adjusted R^2**			
Beta \ mean	**K=1**	**3**	**12**	**K=60**	**K=1**	**3**	**12**	**K=60**	**K=1**	**3**	**12**	**K=60**
L=1	19.16%	17.92%	18.28%	18.07%	-4.26%	-3.12%	-3.63%	-8.39%	21.45%	24.00%	22.36%	17.76%
	(3.66)	(3.71)	(3.74)	(3.40)	(-0.82)	(-0.64)	(-0.72)	(-1.00)				
	[3.67]	[3.76]	[3.76]	[3.59]	[-0.83]	[-0.64]	[-0.74]	[-1.07]				
3	16.17%	15.13%	15.41%	15.28%	-1.88%	-0.83%	-0.94%	-2.35%	24.19%	27.45%	24.94%	20.46%
	(3.81)	(3.93)	(3.89)	(3.15)	(-0.45)	(-0.21)	(-0.23)	(-0.37)				
	[3.79]	[3.94]	[3.96]	[3.31]	[-0.44]	[-0.21]	[-0.26]	[-0.44]				
12	16.77%	15.61%	15.66%	15.59%	-2.45%	-1.27%	-1.23%	-3.07%	21.94%	25.32%	23.31%	18.82%
	(3.82)	(3.93)	(3.88)	(3.24)	(-0.55)	(-0.30)	(-0.28)	(-0.45)				
	[3.64]	[3.50]	[3.83]	[3.19]	[-0.50]	[-0.24]	[-0.28]	[-0.42]				
60	12.90%	13.80%	13.79%	13.31%	0.87%	0.30%	0.65%	0.93%	3.15%	3.52%	4.39%	5.95%
	(2.96)	(3.40)	(3.47)	(2.66)	(0.31)	(0.11)	(0.20)	(0.14)				
	[2.58]	[2.80]	[2.78]	[2.71]	[0.35]	[0.20]	[0.26]	[0.15]				
L=120	6.83%	7.45%	7.60%	6.17%	6.64%	6.44%	6.90%	9.48%	6.05%	7.55%	11.86%	16.45%
	(2.03)	(2.40)	(2.64)	(2.25)	(3.74)	(3.67)	(3.58)	(3.57)				
	[2.92]	[3.27]	[3.45]	[2.18]	[3.05]	[3.07]	[3.17]	[3.39]				

Table 7:
Do horizon betas explain horizon returns for characteristics-sorted portfolios? Further evidence

The table is constructed in the same way as Table 6 and repeats the underlying analysis using the 25 size and momentum sorted portfolios in Panel A and the 25 operating profitability and investment sorted portfolios in Panel B from Ken French's website as test assets. The sample period is the same as in Table 6 except for Panel B that uses returns starting only in July 1963.

	Intercept, $\bar{\lambda}_{K,L,0,t}$				**Slope, $\bar{\lambda}_{K,L,1,t}$**				**Average adjusted R^2**			
Beta \ mean	**K=1**	**3**	**12**	**K=60**	**K=1**	**3**	**12**	**K=60**	**K=1**	**3**	**12**	**K=60**
Panel A: Size and momentum sorted portfolios												
L=1	26.55%	25.65%	24.74%	23.21%	-9.65%	-9.11%	-9.70%	-32.00%	14.95%	16.31%	17.70%	14.06%
	(7.05)	(7.55)	(6.80)	(8.63)	(-2.45)	(-2.47)	(-2.45)	(-5.90)				
	[6.41]	[6.79]	[6.23]	[6.19]	[-2.24]	[-2.27]	[-2.33]	[-3.83]				
3	20.41%	19.46%	19.00%	18.48%	-5.19%	-4.38%	-4.34%	-9.22%	17.05%	19.91%	19.77%	13.18%
	(6.27)	(6.65)	(5.78)	(5.84)	(-1.49)	(-1.32)	(-1.22)	(-2.06)				
	[5.92]	[6.09]	[5.35]	[5.32]	[-1.38]	[-1.22]	[-1.13]	[-1.97]				
12	17.44%	16.49%	15.87%	15.28%	-3.27%	-2.33%	-1.77%	-2.90%	13.90%	16.82%	16.05%	9.49%
	(5.75)	(6.03)	(5.20)	(4.95)	(-0.94)	(-0.69)	(-0.49)	(-0.68)				
	[5.11]	[4.84]	[4.90]	[4.60]	[-0.80]	[-0.51]	[-0.47]	[-0.54]				
60	2.84%	3.36%	3.19%	-1.94%	11.80%	11.64%	12.42%	16.36%	11.58%	11.69%	16.61%	28.58%
	(0.73)	(0.94)	(0.92)	(-0.50)	(4.45)	(4.71)	(4.36)	(4.27)				
	[0.84]	[0.99]	[0.84]	[-0.33]	[2.68]	[2.61]	[2.69]	[3.54]				
L=120	6.05%	6.40%	6.56%	4.07%	7.09%	7.16%	7.66%	11.08%	15.13%	15.72%	23.11%	41.49%
	(1.92)	(2.26)	(2.56)	(2.04)	(5.34)	(5.72)	(5.43)	(4.85)				
	[2.75]	[2.77]	[2.48]	[1.26]	[2.62]	[2.55]	[2.52]	[3.04]				
Panel B: Operating profitability and investment sorted portfolios												
L=1	22.42%	21.32%	20.40%	21.41%	-8.18%	-7.46%	-7.65%	-30.57%	10.06%	12.32%	12.70%	16.32%
	(5.81)	(5.87)	(5.36)	(4.69)	(-2.06)	(-1.92)	(-1.85)	(-2.84)				
	[6.10]	[6.32]	[5.51]	[4.80]	[-2.14]	[-2.02]	[-1.90]	[-2.71]				
3	19.89%	18.99%	18.25%	19.30%	-6.22%	-5.54%	-5.52%	-17.58%	9.93%	12.37%	12.56%	15.29%
	(6.37)	(6.51)	(5.89)	(4.88)	(-1.90)	(-1.73)	(-1.63)	(-2.61)				
	[6.41]	[6.82]	[5.97]	[5.07]	[-1.84]	[-1.77]	[-1.65]	[-2.56]				
12	20.23%	19.33%	18.41%	19.25%	-6.63%	-5.96%	-5.81%	-17.97%	8.88%	11.25%	11.88%	15.51%
	(6.61)	(6.72)	(6.09)	(4.97)	(-2.03)	(-1.85)	(-1.69)	(-2.60)				
	[5.85]	[5.78]	[6.26]	[4.70]	[-1.81]	[-1.64]	[-1.74]	[-2.30]				
60	7.54%	7.67%	7.51%	5.94%	4.75%	4.82%	5.33%	8.00%	1.48%	1.85%	2.92%	7.82%
	(2.82)	(3.23)	(3.34)	(2.63)	(2.65)	(2.78)	(2.79)	(2.59)				
	[2.74]	[2.75]	[2.65]	[1.57]	[1.48]	[1.50]	[1.65]	[1.98]				
L=120	7.28%	7.56%	7.61%	5.68%	4.77%	4.68%	4.96%	7.90%	2.96%	3.33%	6.03%	17.12%
	(2.64)	(3.10)	(3.41)	(3.29)	(4.29)	(4.59)	(4.45)	(4.20)				
	[3.36]	[3.66]	[4.38]	[3.35]	[4.65]	[4.92]	[5.89]	[5.65]				

Table 8:
Conditional distributions for sub-samples

The table shows the percentiles of the cross-sectional distribution of beta as a function of the return horizon, conditional on their one-month value $\beta(1)$. All data is obtained from randomly selecting 4,600 stocks and bootstrapping for each stock a sample of 36,000 monthly returns. Stocks and returns are drawn from all common stocks listed on the NYSE, Nasdaq, and AMEX in the period from 1962 to 2022 that have at least ten years of monthly data available in the respective sample. The bootstrap uses the Politis and Romano (1994) approach with random block length equal to six months. Panel A reports the results for the entire sample, Panel B for the sub-sample up to June 1992, and Panel C for the sub-sample starting in July 1992.

Percentile	**1%**	**5%**	**10%**	**25%**	**50%**	**75%**	**90%**	**95%**	**99%**
Panel A: Whole sample									
1m	0.10	0.36	0.50	0.77	1.09	1.43	1.75	1.99	2.61
3m	0.24	0.44	0.58	0.85	1.17	1.60	1.83	2.21	2.91
12m	0.28	0.50	0.60	0.82	1.20	1.59	1.94	2.28	3.18
5y	0.26	0.56	0.65	0.89	1.49	1.90	2.96	3.29	6.35
Panel B: Sample until June 1992									
1m	0.19	0.41	0.54	0.83	1.12	1.39	1.67	1.85	2.21
3m	0.35	0.53	0.63	0.90	1.24	1.53	1.79	2.07	2.36
12m	0.44	0.56	0.73	0.88	1.26	1.60	1.90	2.27	2.46
5y	0.46	0.55	0.80	0.83	1.45	2.24	2.94	4.51	4.13
Panel C: Sample from July 1992 onward									
1m	0.03	0.28	0.40	0.66	1.03	1.46	1.90	2.19	2.75
3m	0.13	0.37	0.42	0.73	1.11	1.51	2.20	2.36	3.18
12m	0.19	0.36	0.32	0.73	1.26	1.58	2.57	2.50	3.74
5y	0.18	0.34	0.26	0.74	1.87	2.27	4.72	4.76	10.66

Table 9:
Conditional distributions for unlevered returns

The table shows the percentiles of the cross-sectional beta distribution for levered and unlevered returns as a function of the return horizon, conditional on their one-month value $\beta(1)$. All data is obtained from randomly selecting 4,600 stocks and bootstrapping for each stock a sample of 36,000 monthly returns. Stocks and returns are drawn from all common stocks listed on the NYSE, Nasdaq, and AMEX in the period from February 1975 to 2022 that have at least ten years of monthly data. The bootstrap uses the Politis and Romano (1994) approach with random block length equal to six months. Unlevered returns are calculated following Doshi et al. (2019) by scaling excess levered equity returns by $1 - L_{i,t-1}$, where $L_{i,t-1}$ is the ratio of the book value of total liabilities to the sum of the book value of total liabilities and the market value of equity. Panel A reports results for the beta of standard, levered returns using only observations for which unlevered returns are available, and Panels B and C for unlevered returns. The betas in Panels A and B are thereby with respect to standard market returns, whereas Panel C uses the asset-weighted average of unlevered returns.

Percentile	**1%**	**5%**	**10%**	**25%**	**50%**	**75%**	**90%**	**95%**	**99%**
Panel A: Levered returns									
1m	0.04	0.30	0.44	0.71	1.04	1.38	1.73	1.98	2.57
3m	0.16	0.38	0.54	0.78	1.18	1.52	1.89	2.18	2.89
12m	0.21	0.36	0.51	0.81	1.23	1.62	2.00	2.15	3.35
5y	0.33	0.32	0.30	0.81	1.76	2.74	2.60	3.02	9.25
Panel B: Unlevered returns									
1m	0.01	0.05	0.09	0.28	0.56	0.85	1.23	1.49	2.03
3m	0.02	0.05	0.09	0.30	0.63	0.91	1.39	1.63	2.31
12m	0.03	0.04	0.09	0.25	0.59	0.91	1.42	1.73	2.71
5y	0.01	0.02	0.06	0.18	0.56	1.01	1.75	2.54	4.84
Panel C: Unlevered returns with asset-weighted unlevered market return									
1m	0.03	0.11	0.19	0.63	1.19	1.79	2.44	2.90	3.83
3m	0.06	0.12	0.19	0.72	1.33	1.92	2.70	3.16	4.15
12m	0.05	0.10	0.20	0.67	1.36	1.98	3.14	3.48	5.00
5y	0.01	0.09	0.18	0.53	1.35	3.25	4.77	5.15	9.36

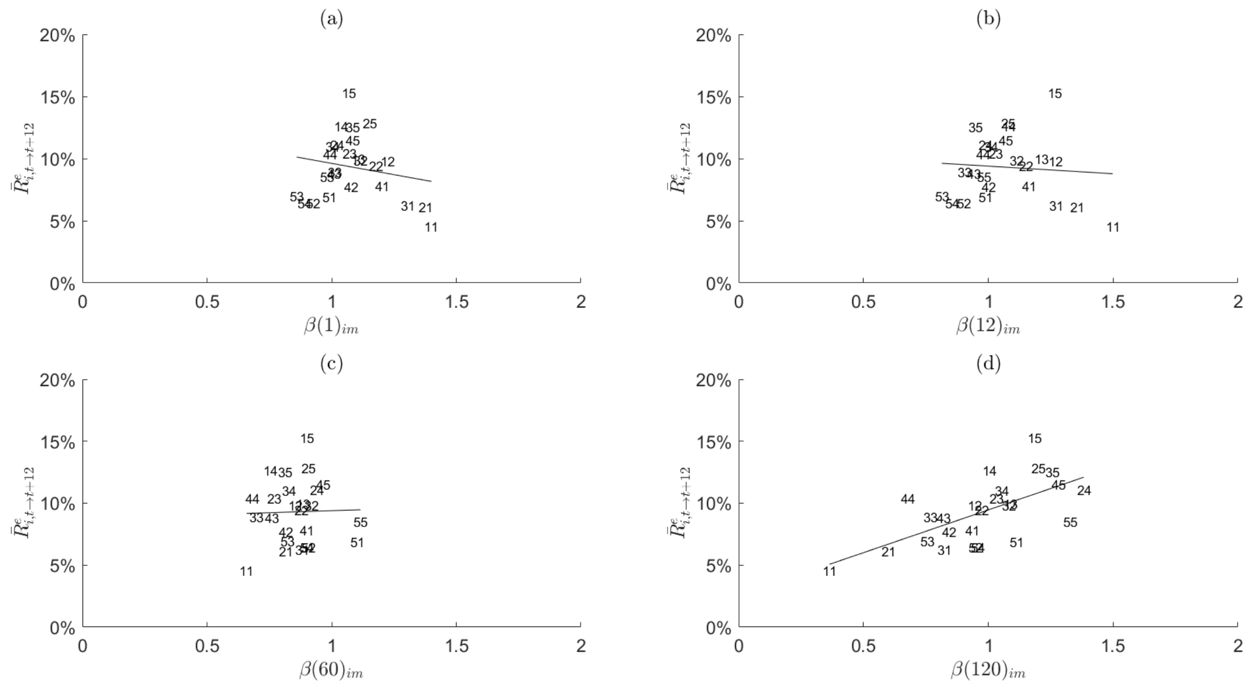

Figure 1: **Annual excess returns and horizon betas for size and B/M sorted portfolios**
The figures plot the average annual returns in excess of the annual risk-free rate and their horizon betas of the 25 size and B/M sorted portfolios. Each two digit number corresponds to a portfolio, whereby the first digit refers to the size quintile (small=1 to big=5) and the second digit refers to the B/M quintile (growth=1 to value=5). The plots also contain a linear trend line in black font. Panels (a), (b), (c), and (d) use horizon betas $\beta(K)$ with K equal to 1, 12, 60, and 120 months, respectively. All portfolio returns are value-weighted. Calculations use overlapping data for horizons beyond one month. The portfolio data is from Ken French's website and the annual risk-free return is from Liu and Wu (2021). The sample period is from January 1962 to December 2022.

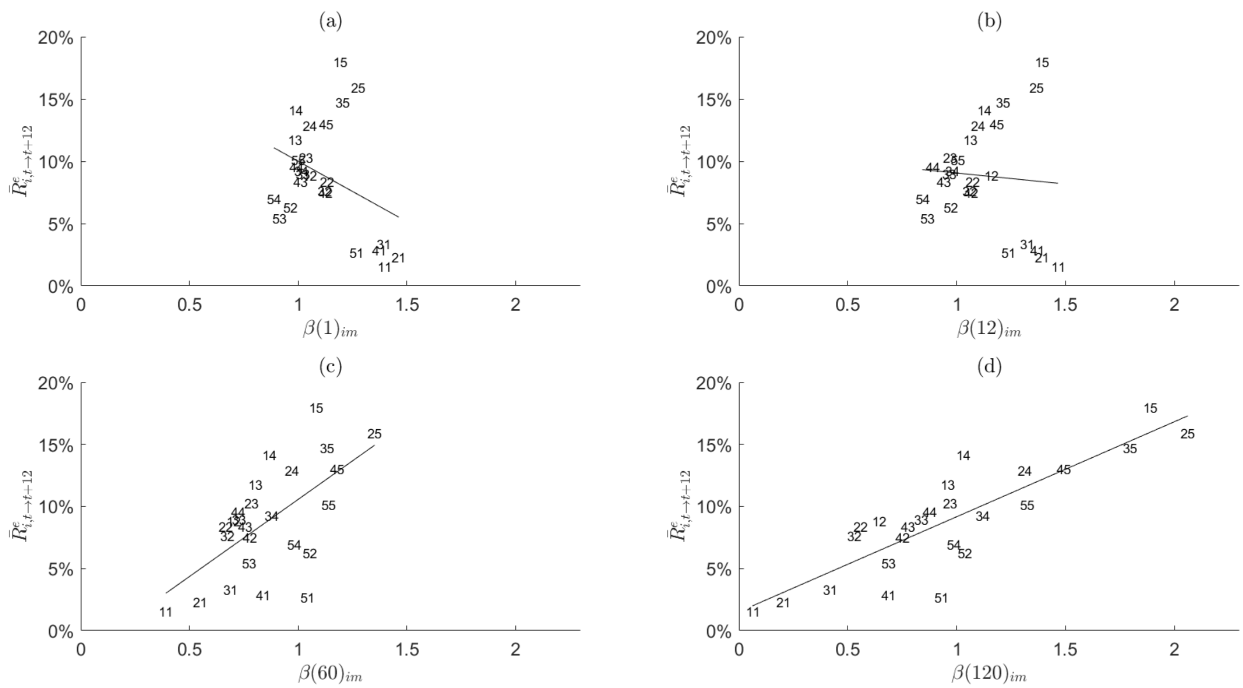

Figure 2: **Annual excess returns and horizon betas for size and momentum sorted portfolios**

The figures plot average annual returns in excess of the annual risk-free rate and their horizon betas of the 25 size and momentum sorted portfolios. The construction of the plots and the underlying calculations and data sources are the same as in Figure 1. The first digit refers to the size quintile and the second digit to the momentum quintile (loser=1 to winner=5). The sample period is from January 1962 to December 2022.

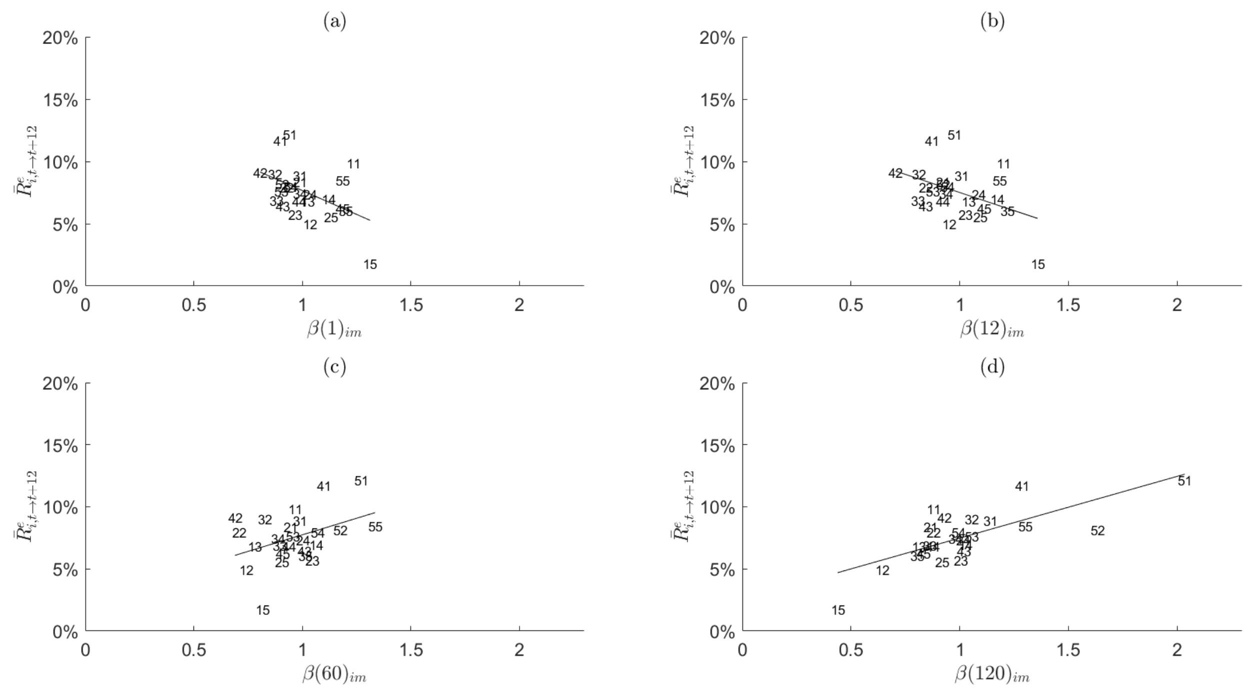

Figure 3: **Annual excess returns and horizon betas for operating profitability and investment sorted portfolios**

The figures plot the average annual returns in excess of the annual risk-free rate and their horizon betas of the 25 operating profitability and investment sorted portfolios. The construction of the plots and the underlying calculations and data sources are the same as in Figure 1. The first digit refers to the operating profitability quintile (weak=1 and robust=5) and the second digit to the investment (conservative=1 to aggressive=5). The sample period is from July 1963 to December 2022.