The Term Structure of CAPM Alphas and Betas

*Wayne Chang**

USC Marshall

Draft: November 18, 2016

— JOB MARKET PAPER —

Abstract

Using monthly returns to estimate portfolio alphas and betas is inappropriate for investors with longer horizons. Alphas and betas have flat term structures only under special conditions that do not hold generally. The paper develops a novel conditional moment estimation method that is simple, non-parametric, and modifies the realized volatility approach to work for long-horizon returns. Long-short portfolios sorted on size, value, and momentum have CAPM betas that can reverse sign with longer horizons. Alphas change too. At multi-year horizons, the average alpha associated with size increases while momentum's decrease until they are of similar magnitudes.

Keywords: Risk-return, Investment horizon, CAPM, Conditional alpha, Time-varying beta, Size effect, Value premium, Momentum effect

JEL Classification: G11, G12, G17

*I would especially like to thank my dissertation chair, Professor Wayne Ferson, for his many hours of patient guidance and advice. For providing valuable input, I'm also grateful to Christopher Jones, Kenneth Ahern, David Solomon, K.R. Subramanyam, Oguzhan Ozbas, Larry Harris, David Harris, Louis Piccotti and to participants in the 2013 GSBA 790 seminar, 2014 USC Finance PhD Brownbag, 2016 SWFA Finance Conference, 2016 USC Finance Brown Bag, and 2016 FMA Annual Conference. All errors are my own. The USC Provost's Ph.D. Fellowship and the USC Graduate School Final Year Fellowship has provided much appreciated financial support. Comments welcome. For contact information, please visit my website at sites.google.com/site/waynelinchang/home.

The risk and return of owning equities depend on the investment horizon. Systematic risk is measured as covariance with a pricing factor (beta) while an investment's abnormal return is what that pricing factor cannot explain (alpha). Under the Capital Asset Pricing Model, the market is the single pricing factor, leading to CAPM alpha and beta estimates ubiquitous in academic and investment communities. The standard approach uses monthly returns with the implicit assumption that monthly alphas and betas are appropriate for everyone. Unlike bond investments, for example, term structure dynamics are rarely considered. This paper shows, however, that alpha and beta term structures are not usually flat. Buy-and-hold investors with multi-year horizons may face different risk-return characteristics than those implied by monthly returns. In some cases, the investment's beta can reverse sign, serving as a hedge instead of adding risk. Using monthly risk-return estimates for all investors is thus inappropriate.

"Horizon," as used in this paper, refers to an investment strategy's anticipated period of implementation. I assume this horizon is exogenous, perhaps due to the investor's preferences or perhaps money is set aside to fund a specific future expenditure. The investment strategy need not be passive but can involve active rebalancing and other preset rules, such as signal-dependent trades, stop-loss rules, or even volatility management. The key requirement is that alpha and beta are estimated using the strategy's past returns and the same strategy remains over the anticipated investment period. The paper explores horizons ranging from 1-month to 10-years. Stock holding periods vary across investors, and these periods can be lengthy. Cella, Ellul, and Giannetti (2013) document for institutional investors a 5%-to-95% range of implied stock holding periods that span 2-months to 6-years. Ameriks and Zeldes (2004) study defined contribution retirement accounts and find that asset allocation changes over the course of a 10-year period occur in less than 30% of them.

The paper studies size, value, and momentum long-short portfolios that do not have fixed holdings but are periodically rebalanced according to the standard Fama and French (1993) approach. Rebalancing ensures the strategies remain exposed to the desired characteristics, even over many years. The paper focuses on CAPM alphas and betas for these characteristics because of their central importance in the asset pricing literature. I study the single market source of systematic risk to simplify the analysis of drivers of sloped alpha and beta term structures. However, the paper's message and methods are applicable more generally. Alphas and betas that reflect multi-factor models can also be estimated, and they can be done for an individual stock, an industry portfolio, or other dynamic portfolios. The paper does not investigate general equilibrium implications or the perspective of the representative investor. It focuses on the empirical documentation of term structure shapes without taking a stand on the economic forces that drive these patterns.

Sloped alpha and beta term structures occur for two theoretical reasons. First, market and individual portfolios can exhibit significant lead-lag correlations with each other. Long-horizon betas aggregate short-horizon contemporaneous betas plus short-horizon lead-lag correlations. Negative correlations produce long-horizon betas smaller than short-horizon ones, resulting in and downward-sloping term structure. I observe this for size when horizons exceed one year. On the other hand, positive lead-lag correlations result in an upward-sloping term structure. I observe this for momentum. Second, alphas and betas are time-varying and mean-reverting. Even without lead-lag correlations, an abnormally high short-horizon beta will likely be followed by less extreme betas, resulting in a long-horizon estimate that's a moderate average. Like interest rates, alpha and beta term structures will tend to be downward-sloping given abnormally high short-horizon estimates and upward-sloping given abnormally low ones. For these two reasons, short-horizon alpha and beta estimates can indeed differ from their long-horizon counterparts. Alphas and betas should thus properly reflect the investment strategy's horizon just like future cashflows should use discount rates with the appropriate maturity. Indeed, betas are inputs into computing discount rates so a wrong beta can lead to misvaluation.

The paper offers two novel contributions. First, it develops a new conditional beta estimation method that uses returns observed at a higher frequency than the investment horizon of interest. For example, to estimate a one-year conditional beta, the method can use daily or monthly returns. This approach is non-parametric, avoiding the use of instruments to capture conditionality (e.g. Ferson and Schadt (1996), Ferson and Harvey (1999)) and of models to describe return behavior. Using high-frequency data is standard in the realized volatility literature (e.g. Andersen and Bollerslev (1998), Fleming, Kirby, and Ostdiek (2003), Barndorff-Nielsen and Shephard (2004), Andersen et al. (2006)). However, the realized approach ignores the conditional mean and estimates quadratic variation rather than the conditional variance needed for betas. It also assumes no lead-lag correlations among high-frequency returns, and thus cannot produce sloped average term structures by construction. Compared to the classical rolling-window approach, the new conditional beta method better captures period-to-period variation and better estimates long-horizon moments that have few adjacent periods.

The paper's second contribution is the documentation of several novel empirical findings regarding long-short portfolios sorted on size (SMB), value (HML), and momentum (UMD). First, it finds significant market-portfolio lead-lag correlations at monthly and annual horizons. They show in-sample and in some cases, out-of-sample predictability. Second, the three portfolios exhibit different average term structure shapes. Size's beta term structure mostly slopes downward, momentum's exhibits the opposite shape, and value's stays mostly flat. Beta term structure effects directly affect alpha term structures since conditional mean return estimates remain similar across horizons. In addition, the covariance term between beta and the market premium has a flat term structure too. Finally, short-horizon conditional alphas and betas are stationary and mean-revert. Consequently, the term structure at a particular point in time can still be sloped even without lead-lag correlations. Overall, these empirical findings have important implications for the estimation and use of alphas and betas, pointing to a more nuanced approach that accounts for the investment horizon. This impacts investor asset allocation and corporate investment decisions requiring beta-dependent discount rates.

Figure 1 previews the paper's main results and shows the average conditional term structure estimated using the new high-frequency approach. Size's beta peaks at the 1-year horizon since monthly lead-lag correlations with the market tend to be positive. With a further lengthening of the horizon, beta's slope reverses and turns downward as negative correlations dominate at annual horizons. When size's beta falls, market comovement explains less of its excess returns. Unexplained alpha thus increases from an insignificant 0.6% at the annual horizon to a significant 5.0% per year at the 10-year horizon. Value and momentum do not share these term structure shapes because they exhibit different lead-lag correlation patterns. Value tends to have insignificant correlations and thus flat alpha and beta term structures. On the other hand, momentum has significantly positive correlations at annual horizons. This produces an upward-sloping beta term structure and alphas that fall from a significant 7.2% at the annual horizon to an insignificant 4.5% per year at the 10-year horizon. I assess the impact of these term structure changes using unconditional

Figure 1: Average Alpha and Beta Conditional Term Structures (New High-Freq Method). Data from 1926-2015 used to form non-overlapping monthly, quarterly, annual returns and annual-overlapping 3-year, 10-year returns. Conditional alpha and beta estimated using the new high-frequency approach described in Section 2. Dotted lines contain the 95% confidence interval calculated using Newey-West '94 standard errors.

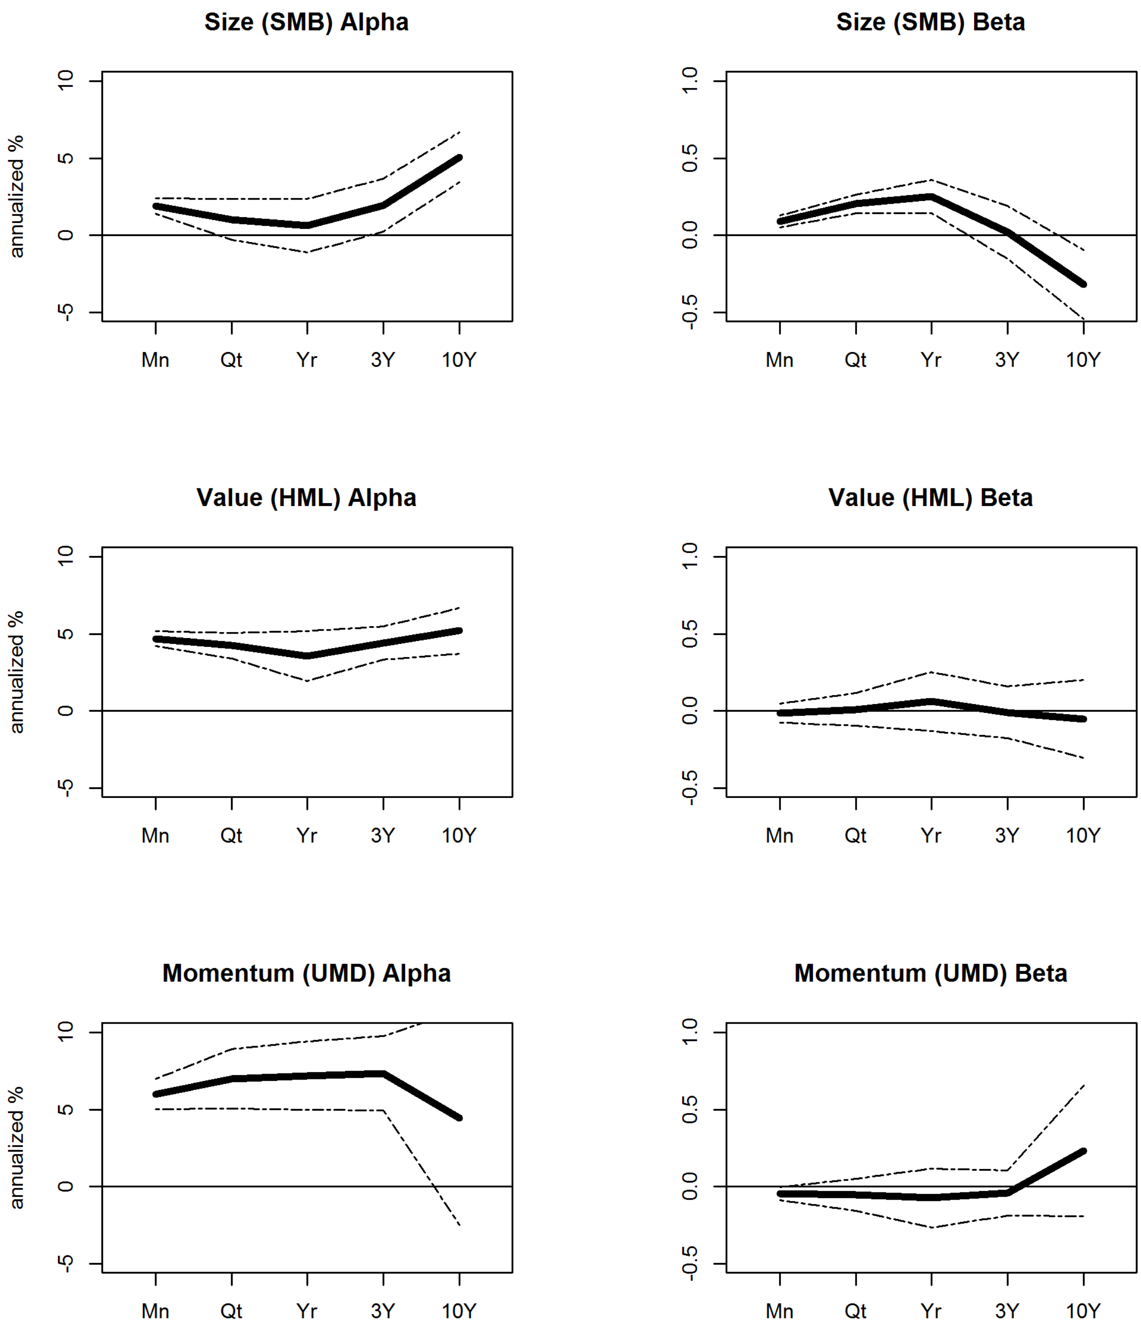

results that assume an investor who seeks a global minimum variance portfolio and has a relative risk aversion of 5. If the investor has a 10-year horizon but improperly optimizes using monthly returns, the certainty-equivalent loss is 9% per year.

This paper builds on the large literature on equity horizon effects. Samuelson (1969) and Merton (1969) demonstrate how the market portfolio's risk-return dynamics depend on the investment horizon. More recent work along this vein includes Campbell and Viceira (2002), Ang and Liu (2004), Campbell and Viceira (2005), Bandi and Perron (2008), Colacito and Engle (2010), Rua and Nunes (2012), Diris, Palm, and Schotman (2014), Chaudhuri and Lo (2015). In particular, the tendency for equity alphas and betas to change with the return horizon is known as the intervalling effect (e.g. Levhari and Levy (1977), Hawawini (1983), K. Cohen et al. (1983), Handa, Kothari, and Wasley (1989), Handa, Kothari, and Wasley (1993), Gencay, Selcuk, and Whitcher (2005), Gilbert et al. (2014)). Past papers usually focus on unconditional alphas and betas and not the time-varying conditional ones I study here. I also analyze a much broader range of horizons, from 1-month to 10-years. For example, Gilbert et al. (2014) find differences between daily and quarterly betas that are mostly induced by positive lead-lag correlations. In contrast, I find the most dramatic effects at multi-year horizons, with *negative* correlations playing the key role.

There's also recent work studying horizon issues of the size, value, and momentum characteristics I examine (e.g. R. Cohen, Polk, and Vuolteenaho (2009), Bandi et al. (2010), In, Kim, and Faff (2010), In, Kim, and Gencay (2011), Jurek and Viceira (2011), Ang and Kristensen (2012), Brennan and Zhang (2013), Kamara et al. (2015)). I differ by investigating the underlying lead-lag correlations that drive alpha and beta horizon effects and by also developing a new conditional estimation approach. I do not explore cross-sectional pricing implications. For example, Kamara et al. (2015) find term structure effects for different factors' cross-sectional price of risk while I study term structure effects for different portfolio's time-series CAPM alpha and beta. Indeed, there's lots of recent work studying term structure dynamics for many other finance concepts. These include historical average returns (Boguth et al. (2016)), market-implied expected returns (Van Binsbergen, Brandt, and Koijen (2012), Ang and Ulrich (2012)), optimal cost-of-capital estimates (Levi and Welch (2016)), investor risk-aversion (Andries, Eisenbach, and Schmalz (2015)), dividend volatility (Belo, Collin-Dufresne, and Goldstein (2015)), tail-risk estimates (Guidolin and Timmermann

(2006)), macroeconomic risks (Boons and Tamoni (2015)), inflation risk price (Ang, Bekaert, and Wei (2008)), variance risk price (Andries et al. (2015)), and consumption risk price (Croce, Lettau, and Ludvigson (2014), Bryzgalova and Julliard (2015), Dew-Becker and Giglio (2016)).

1 Theoretical Sources of Term Structure Effects

This section investigates the drivers of CAPM alpha and beta term-structure dynamics. Alpha estimates follow from betas, so the section begins with an investigation of the determinants of beta term structures. A decomposition of long-horizon betas into short-horizon moments makes clear that beta term structures are only flat under special conditions. The dual presence of lead-lag correlations among short-horizon returns and of mean-reversion in beta estimates induces sloped term structures. Given horizon-dependent betas, alphas are usually horizon-dependent too.

1.1 Decomposing Long-Horizon Beta Into Short-Horizon Betas*

A long-horizon beta estimate can be decomposed into short-horizon market autocorrelation and market-portfolio lead-lag terms. Specifically, the conditional market beta estimate given information at time t and log returns of portfolio i spanning a period of horizon h can be expressed as follows.

$$\begin{aligned}
\hat{\beta}_{t\to t+h} &\equiv \frac{\hat{\mathrm{Cov}}_t(r^i_{t\to t+h}, r^m_{t\to t+h})}{\hat{\mathrm{Var}}_t(r^m_{t\to t+h})} = \frac{\hat{\mathrm{Cov}}_t(\sum_{\tau=1}^{h} r^i_{t+\tau-1\to t+\tau}, \sum_{\upsilon=1}^{h} r^m_{t+\upsilon-1\to t+\upsilon})}{\hat{\mathrm{Cov}}_t(\sum_{\tau=1}^{h} r^m_{t+\tau-1\to t+\tau}, \sum_{\upsilon=1}^{h} r^m_{t+\upsilon-1\to t+\upsilon})} \\
&= \frac{(\hat{\rho}^{im}_{t,h,lag=-h+1} + \dots + \hat{\rho}^{im}_{t,h,lag=0} + \dots + \hat{\rho}^{im}_{t,h,lag=h-1})(\hat{\sigma}^{i}_{t,h,lag=0}\hat{\sigma}^{m}_{t,h,lag=0})}{(\hat{\rho}^{m}_{t,h,lag=-h+1} + \dots + \hat{\rho}^{m}_{t,h,lag=0} + \dots + \hat{\rho}^{m}_{t,h,lag=h-1})(\hat{\sigma}^{m2}_{t,h,lag=0})} \\
&= \frac{(\hat{\rho}^{im}_{t,h,lag=0} + \sum_{l=-h+1}^{h-1,l\neq 0} \hat{\rho}^{im}_{t,h,lag=l})(\frac{\hat{\sigma}^{i}_{t,h,lag=0}}{\hat{\sigma}^{m}_{t,h,lag=0}})}{1 + 2\sum_{l=1}^{h-1} \hat{\rho}^{m}_{t,h,lag=l}} \qquad (1) \\
&= \frac{\hat{\beta}_{t,h,l=0} + \sum_{l=-h+1}^{h-1,l\neq 0} \hat{\beta}_{t,h,lag=l}}{1 + 2\sum_{l=1}^{h-1} \hat{\rho}^{m}_{t,h,lag=l}} \qquad (2)
\end{aligned}$$

The hat " ˆ " symbol denotes sample estimates. Without loss of generality, I assume the short-horizon return spans a period of 1 unit while the long-horizon one spans h units, where h is an integer larger than 1. In line 1, I can thus disaggregate period h returns into the higher-frequency returns indexed by τ and υ. I then assume for line 2 that returns are conditionally covariance stationary, permitting me to consistently estimate the h-horizon covariance terms as a series of higher-frequency correlations.[1] The numerator consists of lead-lag cross-correlations ($\hat{\rho}^{im}_{t,h,lag=l}$) while the denominator consists of market autocorrelations ($\hat{\rho}^{m}_{t,h,lag=0}$). The σ scaling terms are the high-frequency portfolio and market standard deviations. Subscripts t and h mean these terms are estimated using information conditional at t for the period spanning horizon h while the numeric index l denotes the number of high-frequency leads or lags. Equation (1) simply groups together all non-contemporaneous lead-lag correlations. A long-horizon beta equals the properly-scaled sum of higher-frequency market-portfolio cross-correlations divided by the sum of higher-frequency market autocorrelations. Equation (2) converts correlations into betas. The result is a more general case of the classic Scholes and Williams (1977) non-synchronous beta, which adjusts betas for correlated lead-lag returns induced by micro-structure noise. As this derivation shows, Scholes-Williams betas estimate a longer-horizon beta by explicitly incorporating particular choices of lagged betas and autocorrelations. Dimson (1979) betas are also used to correct for micro-structure noise but its definition differs from Scholes-Williams and equation (2) since Dimson betas suffer from inconsistency (Fowler and Rorke (1983)).

The long-horizon beta equals the short-horizon zero-lag beta (ie. $\hat{\beta}_{t\to t+h} = \hat{\beta}_{t,h,lag=0}$) *only if* all numerator and all denominator lead-lag correlations cancel. This occurs, for example, if all lead-lag correlations equal zero. Furthermore, only if this special case holds across all possible horizons is the *average* beta term structure flat. I emphasize average because $\hat{\beta}_{t,h,lag=0}$ can be viewed as the average short-horizon beta estimated for the period spanning t to $t+h$. Therefore, the first driver of sloped beta term structures is the significant presence of market autocorrelations or market-portfolio lead-lag correlations. This idea is old and dates back to at least Levhari and Levy (1977). Lead-lag cross-correlations in the numerator shift the long-horizon beta away from the short-horizon one to the point where the sign can even reverse, as is the case for size and momentum. On the other hand,

[1] Despite having more covariance terms, closer lead-lag correlations are not weighted more heavily in equations (1) or (2). The definition of lead-lag correlation estimates already accounts for the different number of product terms, $\hat{\rho}^{im}_{t,h,lag=l} \equiv \frac{\sum_{\tau=1}^{h-l}(r^i_{\tau+l} r^m_\tau)}{\sqrt{\sum_{\tau=1}^{h}(r^{i2}_\tau)}\sqrt{\sum_{\tau=1}^{h}(r^{m2}_\tau)}}$.

market autocorrelations in the denominator preserve the sign but modulate the magnitude of the long-horizon beta. Empirically, this denominator effect exhibits more limited impact.

In addition to the presence of lead-lag correlations, mean-reverting short-horizon betas are the second driver of sloped beta term-structures. Notice that the short-horizon lag-0 beta discussed above, $\hat{\beta}_{t,h,lag=0}$, may not equal the short-horizon conditional beta estimated at time t, $\hat{\beta}_{t\rightarrow t+1}$. Both are of the same frequency with a short-horizon of 1 unit, and both are conditional on information at time t. But the former is an unconditional estimate of all zero-lag short-horizon betas from time t to $t+h$ while the latter concerns only the first period following time t. The two estimates are similar only if short-horizon betas remain constant over period h or if betas have unit roots. Unit root betas mean further away beta estimates are martingales and thus have an expectation equal to the same horizon beta next period. Therefore, even without auto or lead-lag cross-correlations (such that $\hat{\beta}_{t\rightarrow t+h} = \hat{\beta}_{t,h,lag=0}$), the time t conditional long-horizon beta ($\hat{\beta}_{t\rightarrow t+h}$) differs from the time t conditional short-horizon beta ($\hat{\beta}_{t\rightarrow t+1}$) if short-horizon betas are time-varying and stationary. Indeed, the literature does view betas as stochastic and mean-reverting processes (e.g. Andersen et al. (2006)). Abnormally high betas tend to be followed by lower ones, such that a high $\hat{\beta}_{t\rightarrow t+1}$ will tend to be followed by a lower $\hat{\beta}_{t+1\rightarrow t+2}$. The beta term structure will thus be downward-sloping since $\hat{\beta}_{t\rightarrow t+h}$ is then likely lower than $\hat{\beta}_{t\rightarrow t+1}$. The opposite occurs given abnormally low betas, which tend to be followed by higher betas and a upward-sloping term structure.

Understanding the difference between the term structure at a specific point in time and the average term structure over all time helps differentiate between the two term structure drivers. The first driver, auto and cross-correlations, can produce both date-specific and average sloped term structures. A non-zero $\hat{\rho}^{im}_{t,h,lag=1}$ affects the time t term structure while a non-zero $\hat{\mathrm{E}}[\rho^{im}_{t,h,lag=1}]$ affects the average term structure. The latter is a much stronger condition, requiring multiple time t $\hat{\rho}^{im}_{t,h,lag=1}$ in the historical record to push consistently in the same direction. On the other hand, the second driver, beta mean-reversion, can produce date-specific sloped term structures but not average ones. Without lead-lag correlations, mean-reversion inducing term structure effects will be averaged out across time because $\hat{\mathrm{E}}[\beta_{t\rightarrow t+1}] \approx \hat{\mathrm{E}}[\beta_{t,h,lag=0}]$. The two unconditional average betas are similar because both are the same horizon, estimate contemporaneous comovement, and use the same historical data. Section 4 analyzes the *historical average* term structure where only *historical* *average* auto and lead-lag correlations can explain their slopes. Section 5 assumes away all lead-lag correlations and analyzes how beta mean-reversion induces *date-specific* sloped term structures.

1.2 How Beta Determines Alpha

Given a conditional beta estimate and conditional mean returns, the conditional alpha estimate follows.

$$\hat{\alpha}_{t\to t+h} = \hat{\mathrm{E}}_t[r^i_{t\to t+h}] - \hat{\beta}_{t\to t+h}\hat{\mathrm{E}}_t[r^m_{t\to t+h}] \quad (3)$$

$$\hat{\mathrm{E}}[\alpha_{t\to t+h}] = \hat{\mathrm{E}}[r^i_{t\to t+h}] - \hat{\mathrm{E}}[\beta_{t\to t+h}]\hat{\mathrm{E}}[r^m_{t\to t+h}] - \hat{\mathrm{Cov}}[\beta_{t\to h}, \hat{\mathrm{E}}_t[r^m_{t\to t+h}]] \quad (4)$$

Equation (3) expresses the date-specific conditional alpha estimate while equation (4) expresses the historical average estimate. Jagannathan and Wang (1996) shows how the covariance term can play a critical role. The average conditional alpha term structure is thus a result of the term structures of its four components. The paper's use of log returns produces flat expected mean return term structures. Section 4 estimates a flat term structure for the covariance between betas and market premiums. This leaves the average beta term structure driving the average alpha term structure.

2 Data and Methods Used For Estimating Term Structures

This section covers the data and the empirical methods for estimating CAPM alphas and betas. For each portfolio, three different term structures are estimated: an unconditional, a rolling-window conditional, and a new high-frequency conditional approach.

2.1 Data Source

The paper uses monthly data from Ken French's website and collects market, size (SMB), value (HML), and momentum (UMD) excess-return portfolios for the period from 1927 through 2015. Monthly returns are used to avoid microstructure issues that may impact higher-frequency daily or weekly returns.[2] Gross returns are converted to log returns and compounded by adding them.[3] To get a sense of the term structure, I study monthly, quarterly, annual, 3-year, and 10-year horizons. Multi-year results use overlapping annual returns to increase power. For shorter monthly and quarterly horizons, the high-frequency conditional beta estimation approach necessitates the use of daily returns. Intra-daily data are not used because they have relatively short histories.

All empirical analyses was done with R (RStudio and Microsoft R Open), and I would like to grateful acknowledge the authors of the packages I used: sandwich (Zeileis (2004)), gmm (Chausse (2010)), data.table (Dowle et al. (2015)), zoo (Zeileis and Grothendieck (2005)), dplyr (Wickham and Francois (2016)), magrittr (Bache and Wickham (2014)), tidyr (Wickham (2016)), rmarkdown (Allaire et al. (2016)).

2.2 Unconditional Alpha and Beta

Unconditional alphas and betas simply use OLS regression of portfolio excess returns on market excess returns. Each horizon h requires a separate regression.

$$r^i_{t\to t+h} = \hat{\alpha}^U_h + \hat{\beta}^U_h r^m_{t\to t+h} \quad \forall t = 1, h+1, 2h+1, ..., T-h \tag{5}$$

T denotes the total number of monthly observations and equals 1,068 in this paper (1927 through

[2] Higher-frequency returns may not reflect actual market prices, but admittedly, the distinction between noise and true prices isn't clear-cut. W. Liu and Strong (2008) and Asparouhova, Bessembinder, and Kalcheva (2013) argue that even monthly returns may be too short since posted prices may not be tradeable in large volumes or with acceptable transaction costs. I analyze much longer horizons precisely because I agree it's unrealistic to assume active monthly monitoring. Nevertheless, to get longer horizon returns, I do compound monthly returns. The paper's appendix shows that multi-year results that compound annual instead of monthly returns are qualitatively similar.

[3] Directly adding excess log returns is equivalent to directly compounding excess gross returns and implies monthly rebalancing between the portfolio's long and short legs. For example, 2-month horizon returns obtained through compounding of two 1-month excess gross returns is equivalent to implementing a long-short strategy for the first month, rebalancing to equal weights between the long and shorts legs, and then repeating the strategy for the second month. $r^{gr,ex}_{t\to t+2} = [1 + r^{gr,ex}_{t\to t+1}] \times [1 + r^{gr,ex}_{t+1\to t+2}] - 1 = [1 + r^{gr,long}_{t\to t+1} - r^{gr,short}_{t\to t+1}] \times [1 + r^{gr,long}_{t+1\to t+2} - r^{gr,short}_{t+1\to t+2}] - 1 = [1 + (1 + r^{gr,long}_{t\to t+1}) - (1 + r^{gr,short}_{t\to t+1})] \times [1 + (1 + r^{gr,long}_{t+1\to t+2}) - (1 + r^{gr,short}_{t+1\to t+2})] - 1$. I do not compound long and short legs separately because I analyze size, value, and momentum as dynamic strategies rather than as fixed assets. This approach is consistent with the way Fama-French form SMB, HML, and UMD, with annual portfolio formation for SMB and HML and monthly formation for UMD. Separate compounding would also confound horizon effects since the long-horizon excess return no longer equals the sum of short-horizon returns. The evolution of the long-leg's weight relative to the short-leg's over the horizon will also play a role too.

2015). The five horizons I study translate to h values of 1, 3, 12, 36, and 120 months.

2.3 Conditional Alpha and Beta

Conditional moments are notoriously difficult to estimate since only one observation exists for each time t estimate. In the context of the paper, this reality is problematic for two reasons. First, a variance cannot be estimated with a single observation because it depends on two parameters (i.e. mean and variance). Second, even if a mean parameter were available exogenously, a variance estimate that squares the demeaned single observation, $(r_{t+1} - E_t[r_{t+1}])^2$, is likely too noisy to be useful. Since betas are the ratio of covariance and variance terms, beta estimates further suffer from the noise's non-linear compounding. Meaningful estimates of conditional moments thus require multiple observations. This paper uses two nonparametric methods that avoid an explicit model of returns or return moments. However, the need for multiple observations requires constraints on how conditional moments change. The rolling window approach assumes conditional moments remain constant *across* multiple observations of the return horizon while the high-frequency approach assumes conditional moments remain constant *within* the return horizon.

The rolling-window approach is old and is often used to estimate conditional CAPM alphas and betas (e.g. Fama and MacBeth (1973), Merton (1980), Lewellen and Nagel (2006)). Assuming conditional moments remain constant across time allows an estimate to use adjacent returns of the same horizon. This method is normally used on shorter horizon returns, but the paper implements it on multi-year returns too. Alpha and beta estimates use an OLS regression with the previous J observations of h-horizon returns. Each horizon h and time t requires a separate regression.

$$r^i_{t-\tau\to t+h-\tau} = \hat{\alpha}^{C.Roll}_{t\to t+h} + \hat{\beta}^{C.Roll}_{t\to t+h} r^m_{t-\tau\to t+h-\tau} \quad \forall \tau = h, 2h, ..., Jh \tag{6}$$

For the estimates to be consistent, the conditional alpha and beta during the window period Jh need to be constant. Backward looking windows avoid look-ahead bias, and the base set of results uses window lengths of 5, 10, 20, 30, and 40 years for monthly, quarterly, annual, 3-year, and

10-year horizons, respectively. Five years for monthly returns is standard in the literature, and window lengths rise for longer horizon returns under the constraint that only 89 years of data exists. This constraint means longer horizon returns have windows with fewer observations (i.e. J = 60, 40, 20, 10, and 4, respectively) even as they are longer in duration. The Appendix explores alternative specifications for robustness, and average term structure results remain largely unchanged. Regardless of the window length, average results require an unconditional mean over all the data. Different window lengths may thus affect the persistence and volatility of moment estimates but less so the average level.

2.4 A New High-Frequency Conditional Estimation Method

The paper develops a novel way to estimate conditional moments using returns of a higher frequency than the horizon of interest. Assuming high-frequency conditional moments remain constant within the horizon allows these estimates to be unbiased. The method modifies the standard realized approach since realized volatility may fail to capture multi-month or multi-year variances. Realized volatility estimates quadratic variation and not conditional variance, and the two concepts can diverge at long horizons. This difference can be seen using a simple decomposition of returns under a continuous-time model without jumps.[4]

$$\begin{aligned} r_{t\to t+h} &= \int_t^{t+h} \mu_s ds + \int_t^{t+h} \sigma_s dW_s \\ &\equiv \mu_{t\to t+h} + M_{t\to t+h} \end{aligned}$$

μ_s is the continuous-time expected return process, σ_s is the instantaneous conditional volatility, and W_s is standard Brownian motion. The sum of all the innovations over horizon h are aggregated into the martingale $M_{t\to t+h}$.

Quadratic variation is the realized square of the martingale innovation, $M^2_{t\to t+h}$. In the absence of jumps, it equals the integrated volatility $\int_t^{t+h} \sigma_s^2 ds$ because Brownian innovations are uncorrelated,

[4]This setup follows Andersen, Bollerslev, and Diebold (2010). They use "expected volatility" to refer to conditional variance and "notional volatility" to refer to quadratic variation.

$\mathrm{Cov}[dW_s dW_v] = 0 \quad \forall s \neq v$. The expected return process is usually considered to have bounded variation so adds no quadratic variation. In contrast, the discrete horizon conditional variance is defined as $\mathrm{Var}_t[r_{t\to t+h}] \equiv \mathrm{E}_t[(r_{t\to t+h} - \mathrm{E}_t[r_{t\to t+h}])^2]$, with $\mathrm{E}_t[r_{t\to t+h}] = \mathrm{E}_t[\mu_{t\to t+h}]$ as the conditional mean over horizon h. Expected quadratic variation, $\mathrm{E}_t[M^2_{t\to t+h}]$, is a component of conditional variance as the following decomposition shows.

$$\begin{aligned}
\mathrm{Var}_t[r_{t\to t+h}] &\equiv \mathrm{E}_t[(r_{t\to t+h} - \mathrm{E}_t[r_{t\to t+h}])^2] \\
&= \mathrm{E}_t[((\mu_{t\to t+h} - \mathrm{E}_t[r_{t\to t+h}]) + M_{t\to t+h})^2] \\
&= \mathrm{Var}_t[\mu_{t\to t+h}] + 2\mathrm{Cov}_t[\mu_{t\to t+h}, M_{t\to t+h}] + \mathrm{E}_t[M^2_{t\to t+h}] \qquad (7)
\end{aligned}$$

$\mathrm{E}_t[\cdot]$ denotes the conditional expectation using information available at time t. Standard realized measures that estimate quadratic variation deliver the third term of equation (7). Thus, they are adequate approximations of conditional variance if the first two terms of (7) are negligible. The approximation works given a sufficiently short horizon when the magnitude of unexpected innovations dominates that of expected returns. In the limit as the horizon shrinks to zero, conditional variance approaches expected quadratic variation, with both equal to the instantaneous variance σ^2_s. The mean return process contributes no variance or covariance, and the $\mathrm{Var}_t[\mu_{t\to t+h}]$ and $\mathrm{Cov}_t[\mu_{t\to t+h}, M_{t\to t+h}]$ terms can be ignored.

With a lengthening of the horizon, the difference between quadratic variation and conditional variance increase as the relative contribution of the expected return component grows. Andersen, Bollerslev, and Diebold (2010) provide a numerical example based on the mean-reverting Ornstein-Uhlenbeck process. They estimate a near-zero deviation between quadratic variation and conditional variance at the daily horizon but a few percentage point difference by the quarterly horizon. The next section shows that by the annual horizon, the relative importance of expected versus unexpected returns can reverse. Lead-lag covariances can dominate contemporaneous covariance such that quadratic variation is no longer a meaningful approximation of conditional variance. At long horizons, the $\mathrm{Cov}_t[\mu_{t\to t+h}, M_{t\to t+h}]$ terms cannot be ignored.

Figure 2 provides a useful comparison of realized volatility and this paper's new variance estimation method. Panel (a) shows that even if lead-lag cross-correlations generate a sloped variance term structure, realized volatility will always estimate a flat average term structure by construction. This flaw extends to betas and is apparent in equation (2). The realized approach ignores lead-lag correlations so even if they exist, a flat term structure ensues because $\hat{\beta}_{t\to t+h} = \hat{\beta}_{t,h,lag=0}$ regardless of h.

Two modifications to the realized approach lead to the new high-frequency estimation method and are necessary for estimating conditional variance instead of quadratic variation.

1. Assume a conditional expected value for the return horizon, $\bar{\mu}_{t\to t+h}$, that proxies for $\mathrm{E}_t[r_{t\to t+h}] = \mathrm{E}_t[\mu_{t\to t+h}]$. The conditional mean process μ_s can be stochastic and be left unspecified but its expected value over horizon h is a constant and must be assumed. In my base case, I use a rolling-average of past returns to specify $\bar{\mu}_{t\to t+h}$.
2. Provide a weighting method for the lead-lag auto-covariances among the high-frequency returns. In my base case, I use a Gaussian kernel where the weights for lead-lag l of horizon h are $w(l,h) = e^{-\frac{(\pi l/h)^2}{2}}$. This gives the concurrent term a weight of 1 with more distant terms receiving smaller and smaller weights. Closer correlations more likely capture true economic relationships that contribute to the conditional variance while more distant correlations more likely reflect noise. Leads and lags are treated symmetrically. **Figure 2** panels (c)-(d) provides a visual depiction.

The following illustrative example is consistent with **Figure 2** and estimates the time t conditional variance of a 5-period horizon return, $\hat{\mathrm{Var}}_t[r_{t\to t+5}]$. It may be helpful to think of a single period as one month. The realized approach estimates quadratic variation as $\hat{M}^2_{t\to t+5} = \sum_{\tau=1}^{5} r_\tau^2$ but quadratic variation may not adequately approximate conditional variance. Alternatively, a naive unbiased estimate of conditional variance squares the demeaned return: $\hat{\mathrm{Var}}_t[r_{t\to t+5}] = \tilde{r}^2_{t\to t+5}$, where $\tilde{r}_{t\to t+5} \equiv r_{t\to t+5} - \bar{\mu}_{t\to t+5}$. Next, high-frequency returns are used by noting that $\tilde{r}^2_{t\to t+5} = (\sum_{\tau=1}^{5} \tilde{r}_\tau)(\sum_{\upsilon=1}^{5} \tilde{r}_\upsilon)$, where $\tilde{r}_\tau$ and $\tilde{r}_\upsilon$ are demeaned high-frequency returns. I demean through an even split of $\bar{\mu}_{t\to t+5}$ to get $\tilde{r}_\tau = r_\tau - \frac{1}{5}\bar{\mu}_\tau$.[5] The second modification simply weights

[5]The decomposition holds regardless of how $\bar{\mu}_{t\to t+h}$ is allocated among the high-frequency returns. But different allocations can affect individual high-frequency cross-terms. Since the proposed high-frequency method incorporates only a subset of these cross-terms, the resulting conditional variance estimate may be sensitive to how $\bar{\mu}_{t\to t+h}$ is

Figure 2: Realized vs. New High-Frequency Approach. Figure illustrates differences between the standard realized approach shown in panels (a)-(b) and the paper's new high-frequency approach shown in panels (c)-(d). The following is an unbiased conditional covariance estimate under rational expectations: $\hat{\text{Cov}}_t[r^i_{t\to t+h}, r^m_{t\to t+h}] = \tilde{r}^i_{t\to t+h}\tilde{r}^m_{t\to t+h}$, where $\tilde{r}$ is a demeaned return. These estimates are shown as large boxes with numeric side labels counting the periods constituting the h-period horizon. Each interval on an axis depicts one period, with market returns on the x-axis and portfolio returns on the y-axis. Panels (a) and (c) depict a sequence of five 5-period horizons while panels (b) and (d) depict a single 25-period horizon. The standard realized approach captures only high-frequency contemporaneous cross-products, represented by the sequence of diagonal black boxes. In contrast, the new approach weights all lead-lags based on a specified weighting kernel, with darker shades representing larger weights. I use a Gaussian kernel. The realized approach's average covariance term structure (comparing 5-period versus 25-period horizons, with proper scaling) is flat by construction since the two horizons capture the same black squares. This is not true for the new approach. More lead-lags exist for the 25-period horizon, and the squares are also weighted differently. The average term structure will thus be upward-sloping if the sign of the sum of these different weighted terms is positive.

(a) Realized Approach (Horizon = 5)

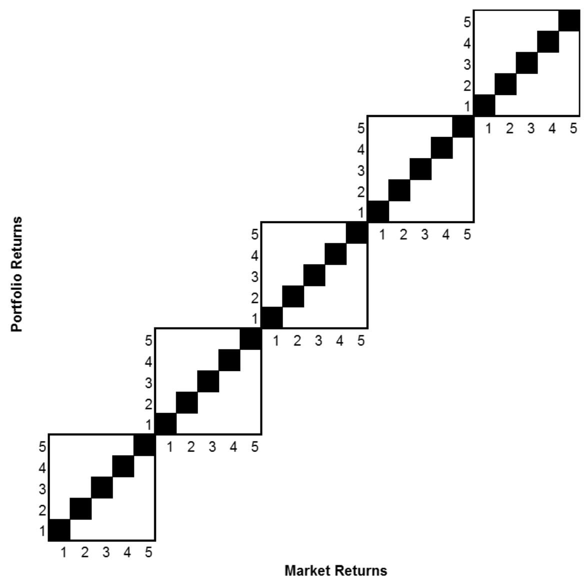

(b) Realized Approach (Horizon = 25)

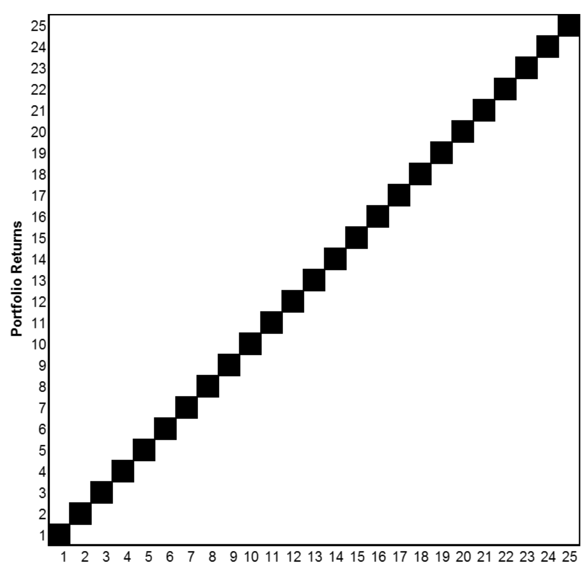

(c) New High-Frequency Approach (Horizon = 5)

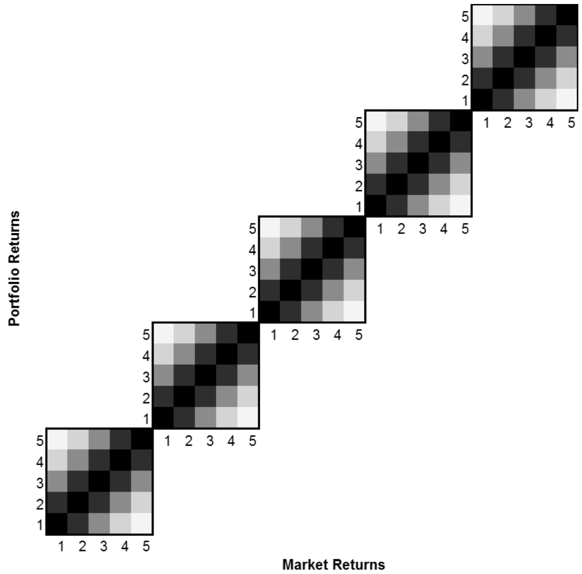

(d) New High-Frequency Approach (Horizon = 25)

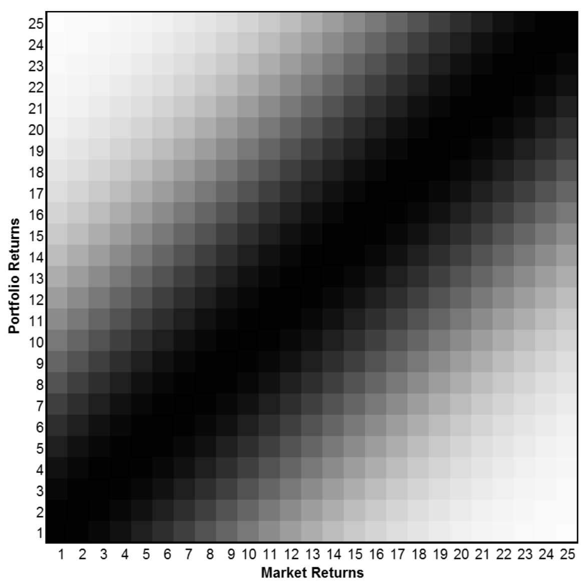

lead-lag correlations based on their distance from contemporaneous comovement. The resulting estimate equals $\hat{\text{Var}}_t[r_{t\to t+5}] = \sum_{\tau=1}^{5} w(0,5)\tilde{r}_\tau^2 + 2\sum_{\tau=1}^{4}(w(1,5)\tilde{r}_{\tau+1}\tilde{r}_\tau) + ... + 2(w(4,5)\tilde{r}_{\tau+4}\tilde{r}_\tau)$ where $w(0,5) = 1$, $w(1,5) \approx 0.82$, $\ldots$ $w(4,5) \approx 0.04$ under the Gaussian kernel.

allocated. I choose an even allocation because it is simplest. I did not investigate alternative allocation strategies and suspect that reasonable alternatives are unlikely to result in qualitatively different estimates.

The standard realized method is a special case of the new high-frequency method, with a conditional mean assumption of $\bar{\mu}_{t \to t+h} = 0$ and only the lag-0 term receives a unit weight. To the extent that these assumptions are appropriate, the realized approach does estimate conditional variance. But as the following section demonstrates, not weighting any lead-lag correlations is likely inappropriate for multi-month and multi-year portfolio returns. The realized approach estimates quadratic variation, but quadratic variation no longer approximates conditional variance well. And it is conditional variance estimates that are needed.

The Scholes-Williams beta can also be considered a special case of the new high-frequency method, where a preselected set of lags is uniformly unit-weighted while all other lead-lags are ignored. Traditionally, Scholes-Williams beta corrects for microstructure noise, so included lags are the ones where these effects spill into. In addition, the conditional mean assumption $\bar{\mu}_{t \to t+h}$ equals the realized return over horizon h because Scholes-Williams betas use standard regressions that effectively assume that. This is problematic, however, since the realized return is now being used to estimate both the mean and the variance. When all lead-lags are equally weighted, this results in a meaningless zero variance estimate. Scholes-Williams normally doesn't weight all lead-lags, but those that are ignored are ignored because they have zero expectation. Therefore, to the extent that the resulting estimate is non-zero, it must be estimating bias among the incorporated lags and not the actual variance. For this reason, the new high-frequency approach first demeans using an exogenously specified mean and then squares (and weights) the resulting demeaned returns. It avoids using a standard regression to estimate lead-lag betas precisely because that implicitly and inappropriately sets the conditional mean assumption to the realized return.

Using the new high-frequency approach to estimate conditional betas reflects the result in equation (2). Alphas directly follow using equation (3). Specifically, I use realized returns at time $t+1$ to estimate conditional expectations formed at time t for the following $t+1$ period. This approach follows the realized literature where time $t+1$ realized variance also estimates conditional variance formed at time t. As long as there are no systematic biases in expectations, the realized observation is an unbiased estimate of the conditional expectation. To address concerns of look-ahead bias, an alternative interpretation uses the time $t+1$ realized estimate as the conditional expectation at $t+1$ for time $t+2$. This difference does not affect the paper's main empirical results since these results average across all conditional expectations.

One wrinkle is that a lower limit, $Min.Den$, needs to be placed on the beta denominator term since the inclusion of large negative autocorrelations may dramatically shrink and even turn the denominator negative. The denominator estimates long-horizon market variance so it should clearly be positive. A positive $Min.Den$ enforces this while a non-epsilon value prevents the denominator from being too small and blowing up the beta estimate. The following summarizes the high-frequency approach for estimating alphas and betas of horizon h using information conditional at time t. Each horizon h and time t requires a separate estimate.

$$\begin{aligned}
\hat{\beta}^{C.HF}_{t\to t+h} &\equiv \frac{\left(\sum_{l=-h+1}^{h-1} w(l,h)\hat{\rho}^{im}_{t,h,lag=l}\right)\left(\frac{\hat{\sigma}^{i}_{t,h,lag=0}}{\hat{\sigma}^{m}_{t,h,lag=0}}\right)}{Max\left(\sum_{l=-h+1}^{h-1} w(l,h)\hat{\rho}^{m}_{t,lag=l}, Min.Den\right)} \\
\hat{\alpha}^{C.HF}_{t\to t+h} &\equiv \bar{\mu}^{i}_{t\to t+h} - \hat{\beta}^{C.HF}_{t\to t+h}\bar{\mu}^{m}_{t\to t+h} \\
where & \\
\tilde{r^{i}_{\tau}} &\equiv r^{i}_{\tau} - \frac{1}{h}\bar{\mu}^{i}_{t\to t+h} \\
\hat{\sigma}^{i}_{t,h,lag=0} &\equiv \sqrt{\sum_{\tau=1}^{h}(\tilde{r}^{i2}_{\tau})} \\
\hat{\rho}^{im}_{t,h,lag=l} &\equiv \frac{\sum_{\tau=1}^{h-l}(\tilde{r}^{i}_{\tau+l}\tilde{r}^{m}_{\tau})}{\hat{\sigma}^{i}_{t,h,lag=0}\hat{\sigma}^{m}_{t,h,lag=0}} \\
w(l,h) &\equiv \text{weighting kernel for lead-lag}\, l
\end{aligned} \tag{8}$$

$\bar{\mu}^{i}_{t\to t+h}$ and $\bar{\mu}^{m}_{t\to t+h}$ must be exogenously specified and not be the realized return within the horizon. I follow the rolling-window approach with the same base case window lengths and average the previous J periods of horizon h returns. I use the Gaussian weighting kernel $w(l,h) = e^{-\frac{(\pi l/h)^2}{2}}$. I use daily returns as the high-frequency return for monthly and quarterly horizons and monthly returns for annual, 3-year, and 10-year horizons. As long as the high-frequency horizon isn't too coarse, the specific choice shouldn't matter much. Looking at **Figure 2** panel (d), a Gaussian weighting kernel means the weights at specific locations of the grid should be largely unaffected by the grid's granularity. Finally, I choose $Min.Den = 0.3$, which constrains the long-horizon market variance to be at least 30% of the scaled short-horizon variance. This limit binds only in rare circumstances when dramatic reversion in market returns occurs during the period. All these specifications are somewhat arbitrary, so the Appendix explores alternatives and finds broadly similar results.

3 The Presence of Auto and Lead-Lag Correlations

Persistent market-portfolio cross-correlations or market autocorrelations are the only way to produce a sloped *average* beta term structure. This possibility is often dismissed since substantive lead-lag correlations contain predictive and potentially profitable information that should be arbitraged away. However, predictability need not imply market inefficiency (Fama and French (1989)), and the literature has consistently documented both market return autocorrelation (e.g. Fama and French (1988), Poterba and Summers (1988), Conrad and Kaul (1989)) and portfolio cross-correlations (e.g. A. Lo and MacKinlay (1990), Chordia and Swaminathan (2000), Hou (2007), Hong, Torous, and Valkanov (2007), L. Cohen and Frazzini (2008), Menzly and Ozbas (2010), Chordia, Sarkar, and Subrahmanyam (2011), L. Cohen and Lou (2012)). Different explanations have been proposed, including time-varying risk aversion, liquidity shocks, differences in the speed of information diffusion, and supply-chain links. This section empirically evaluates the presence of market autocorrelations and market-portfolio cross-correlations in monthly and annual returns. I ignore daily returns because micro-structure noise leading to daily or weekly lead-lags is already well-known and will play only a minor role in the longer-horizon results I examine.

3.1 Significance of Individual Auto and Lead-Lag Correlations

Figure 3 shows standard auto and cross-correlograms using all historical data, with exceedances beyond the dotted lines indicating a 95% significance level. Significance is based on a single test, but the figure shows results for multiple lead-lags. For every 20 lead-lags shown, therefore, we should expect only one exceedance occurring by chance. Left and right panes show monthly and annual correlations, respectively. Panel (a) depicts excess-market autocorrelations that are adjusted for finite-sample bias (Kendall (1954)). Although there is some evidence of significant monthly correlations, no exceedances occur at the annual horizon. What can cause sloped beta term structure are sequential correlations of nearby lags that push in the same direction. Otherwise, if some correlations are negative while others are positive, they cancel and leave no net impact. This directional consistency, however, is not apparent in market returns, suggesting a limited contribution from the market denominator component of term structure effects (see equation (1)).

Panels (b)-(d) show market-portfolio correlations for size, value, and momentum, with positive lags indicating market returns that occur before and thus anticipate portfolio ones. Results here are interesting and foreshadow beta term structure patterns in the next section. Monthly lag-0 correlations can be easily converted into monthly betas (simply scale by $\frac{\sigma^i_{month}}{\sigma^m_{month}}$) while annual lag-0 correlations similarly reflect annual betas. Ignoring market denominator effects, equation (1) shows that long-horizon betas are simply the sum of higher-frequency lead-lag correlations. Visually summing these correlations tells us the direction in which average longer-horizon betas slope. The same scale for monthly and annual results is maintained to facilitate this exercise.

Let's begin with size in panel (b). Monthly exceedances abound, with the 1-month lagged correlation being the most striking. Gilbert et al. (2014) argue that small stock prices incorporate information more slowly because they are likely less-understood and more opaque. It thus takes time for market changes to catch up to small stocks, resulting in rising betas as the horizon lengthens and fundamental co-movement is eventually reflected more fully. Interestingly, this pattern reverses itself as the number of lags increase, with negative correlations becoming more prevalent and having greater significance. Annual results make this pattern apparent and show a sequence of 5 negative correlated lags. The second is especially significant and large enough to counteract the positive lag-0 correlation. This anticipates the next section when we see SMB's historical-average beta term structure peak at the annual horizon and then decline to become eventually negative at multi-year horizons. It's unlikely that gradual information diffusion explains this finding since negative correlations reflect reversal and not return continuation.

The results for value in panel (c) are interesting in a different way. The lag-0 monthly correlation is significantly positive, but the magnitude shrinks by half and becomes insignificant for the lag-0 annual correlation. Indeed, at the annual horizon, non-contemporaneous lead-lag correlations have much larger magnitudes. But these correlations have little significance and mostly cancel. This suggests that estimation noise will dominate at the annual horizon, making it difficult to reject the

Figure 3: Auto and Lead-Lag Correlations at Individual Lead-Lags. Data from 1926-2015 used to form non-overlapping monthly (left panes) and annual returns (right panes). **Panel (a)** shows market excess return autocorrelations at individual lead-lags, with adjustments for finite-sample bias using Kendal '54. **Panels (b)-(d)** show market-portfolio lead-lag correlations. Dotted lines contain the 95% confidence interval under the null of iid returns leading to zero correlation at all lags. This means $\hat{\rho}_l$ is asymptotically normal with mean zero and variance $\frac{1}{T}$.

(a) Market Autocorrelations

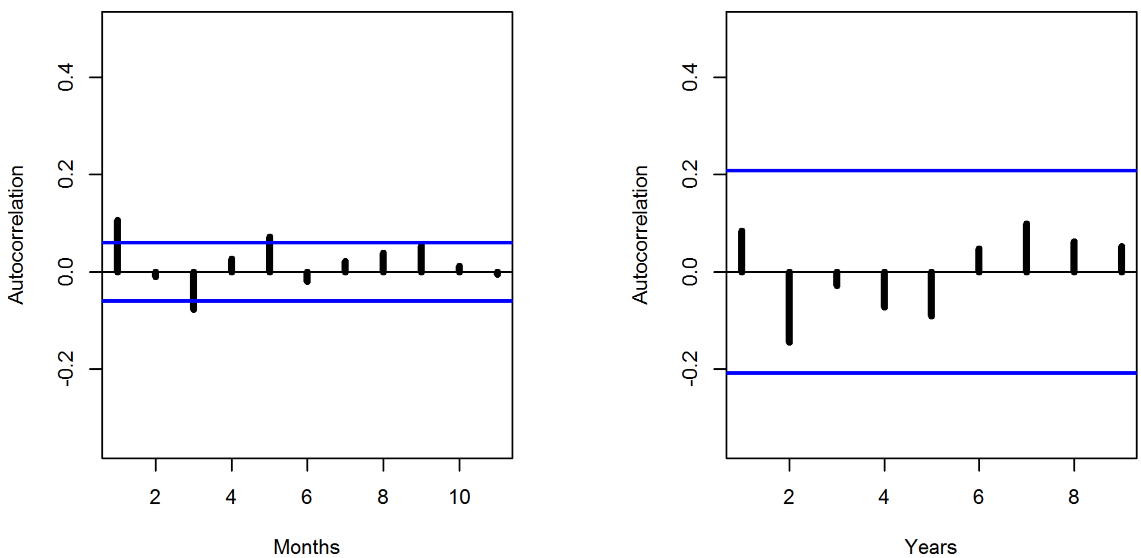

(b) Market-Size Cross-Correlations

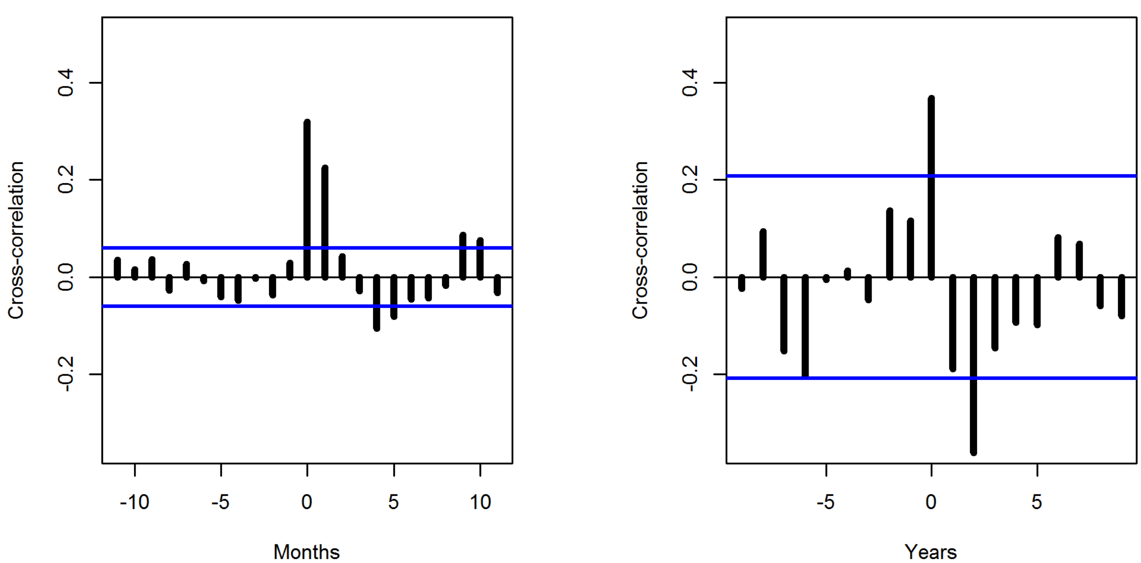

(c) Market-Value Cross-Correlations

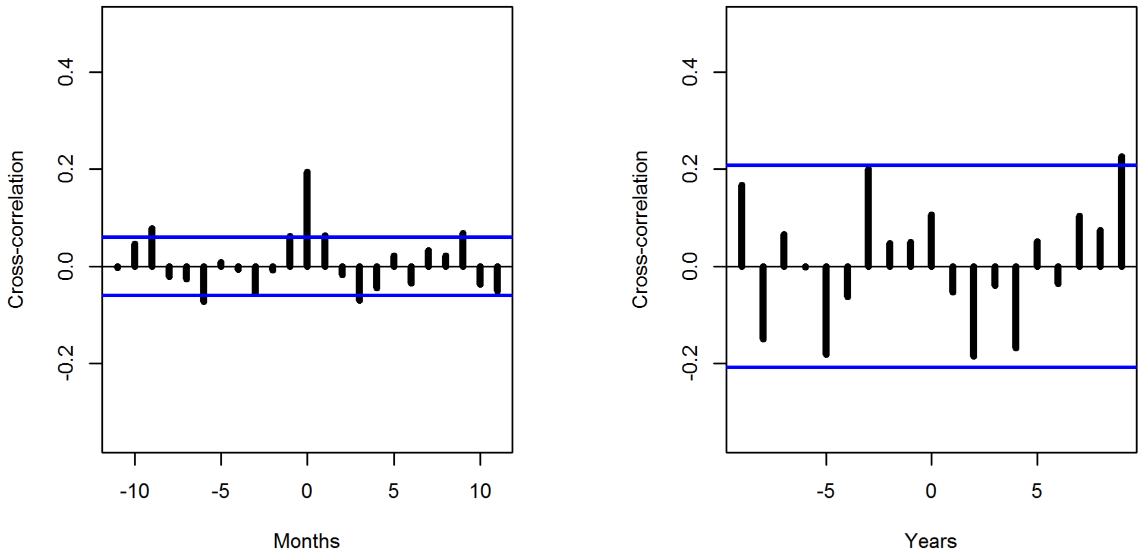

(d) Market-Momentum Cross-Correlations

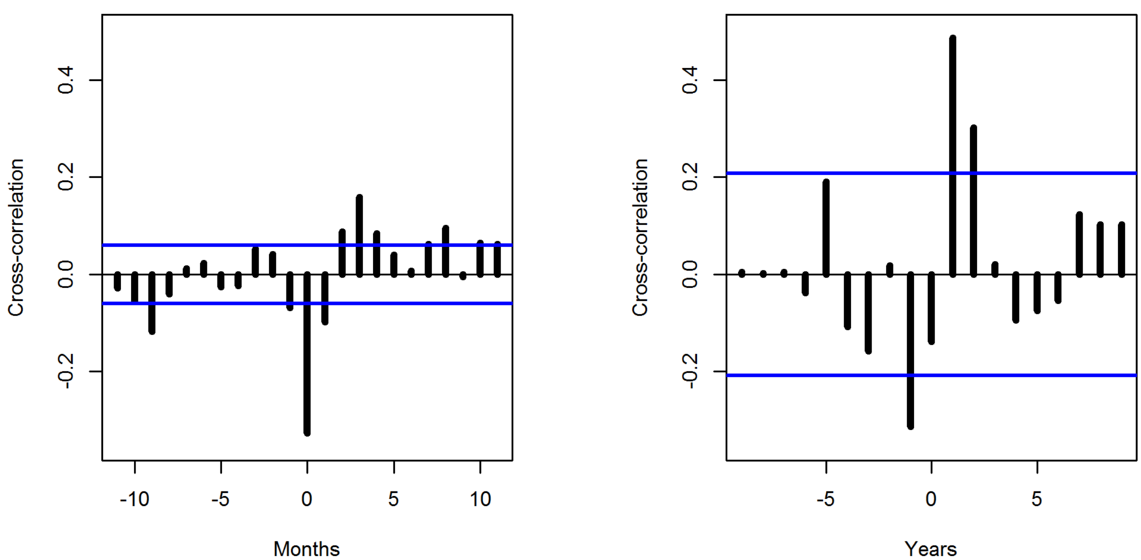

flat beta term structure hypothesis. Momentum's results in panel (d) are like size's but in reverse. At the monthly horizon, the lag-0 correlation is very negative, but significant exceedances at further lead-lags go the other way and tend to be positive. By the annual horizon, the lag-0 correlation loses its significance and is in fact, dwarfed by highly positive lag-1 and lag-2 values. This points to a term structure pattern that's opposite to size, where monthly betas tend to be negative but then slope upward to become positive at longer horizons.

3.2 Significance of Cumulative Auto and Lead-Lag Correlations

Beta term structure effects require adjacent lead-lag correlations to push in the same direction consistently. While **Figure 3** looks at individual lags, **Table 1** sums them cumulatively to capture better when correlations of neighboring lead-lags agree and are more likely to have an impact. Leads and lags of the same distance are summed since their impact on the beta term structure occur together. Panel (a) shows in-sample results that use all historical data, like the **Figure 3** results we already looked at. Market autocorrelations are positive at monthly horizons, in some cases significantly so, but they become negative by Year-2. Cumulative autocorrelations peak at -25% by Year-5, reflecting the mean-reversion tendencies of 3 to 5-year returns found in Fama and French (1988). The market denominator impact is thus to pull betas toward zero at monthly horizons (when the denominator becomes larger) but magnifies them at multi-year horizons (when the denominator shrinks). These results, however, have little to no statistical significance.

For size, we see cumulative positive correlations that are significant up to a lag of 3-months, but the sign then reverses and becomes significantly negative by Year-5. Value shows little significance and only small magnitudes, implying a mostly flat term structure. Momentum sees mostly positive cumulative lead-lag correlations with significance beginning at Month-3 and also showing up in Year-2. Like size, momentum has large annual lagged correlations but where it differs, as **Figure 3** shows, is that its leads offset much of these effects. Size's leads are relatively smaller in magnitude, and thus the lags dominate. This means that multi-year term structure effects will be more muted for momentum, and we see this in lower cumulative annual correlations that are less significant.

Unlike in-sample panel (a) results, **Table 1** panel (b) assesses how investors could have

Table 1: Auto and Lead-Lag Correlations at Cumulative Lead-Lags. Data from 1926-2015 used to form monthly and annual non-overlapping returns. **Panel (a)** assesses the extent of in-sample predictability. Using all historical data, excess-market autocorrelations and market-portfolio lead-lag correlations are calculated. Correlations of the same distance are combined and closer correlations are summed together $\sum_{l=-lag,l\neq 0}^{l=lag} \rho_l$. This indicates the total directional impact on longer horizon returns of all correlations up to the indicated lag. A significant difference from zero at a two-sided 95% level is boldfaced and assessed using $tstat_{lag} = \frac{\sum_{l=-lag,l\neq 0}^{l=lag} \rho_l}{\sqrt{2lag/T}}$. Standard errors are calculated under the null of iid returns leading to zero correlation at all lags. **Panel (b)** assesses the extent of out-of-sample predictability by showing the OOS $R^2_{lag} = 1 - \frac{MSE_{lag,alternative}}{MSE_{lag,null}}$, with $MSE_{lag,forecast} = \sum_{t=11}^{t=T}(\sum_{l=-lag,l\neq 0}^{l=lag} \rho_{t,l,forecast} - \sum_{l=-lag,l\neq 0}^{l=lag} \rho_{t,l,actual})^2$. The first 30 periods form the training period. For monthly and annual correlations, $\rho_{t,l,actual}$ is calculated using the next 12-month and 10-year period, respectively. The null predicts zero correlation at all lags while the alternative predicts they equal their historical value. Predictability is assessed on a one-step-ahead rolling basis with continuously updated forecasts. Boldfaced R^2 values are positive and indicate historical outperformance by the alternative over the null. Market autocorrelations are adjusted for finite-sample bias using Kendal '54.

(a) In-Sample Cumulative Lead-Lag Correlations (%)

	Months (Cum Lags and Leads)											Years (Cum Lags and Leads)								
	1	2	3	4	5	6	7	8	9	10	11	1	2	3	4	5	6	7	8	9
Market Auto	**11**	**10**	2	4	12	10	12	16	**21**	**22**	**22**	8	-6	-9	-16	-25	-20	-11	-4	1
Market-Size Cross	**25**	**26**	**23**	7	-5	-10	-12	-16	-4	5	6	-7	-30	-49	-57	**-67**	**-80**	**-88**	**-84**	**-95**
Market-Value Cross	**12**	10	-3	-8	-5	-16	-15	-15	-0	1	-5	-0	-14	2	-21	-34	-38	-21	-28	11
Market-Momentum Cross	**-17**	-4	**17**	**23**	**25**	**28**	**35**	**41**	**28**	**29**	**33**	17	**49**	35	15	27	18	30	41	51

(b) Out-Of-Sample R^2 of Cumulative Lead-Lag Correlations (%)

	Months (Cum Lags and Leads)											Years (Cum Lags and Leads)								
	1	2	3	4	5	6	7	8	9	10	11	1	2	3	4	5	6	7	8	9
Market Auto	-11	-9	-1	-0	**1**	**5**	**7**	**10**	**10**	**17**	**15**	-46	-34	-28	-16	-20	-23	-20	-18	-14
Market-Size Cross	-13	**1**	**2**	-1	-11	-0	-5	-19	-85	-85	-72	-18	**9**	-11	-80	-617	-214	-512	-313	-117
Market-Value Cross	-25	**6**	**8**	-3	-7	0	-3	-12	-46	-59	-78	-13	-17	-48	-32	-557	-274	-326	-320	-106
Market-Momentum Cross	**3**	-4	-20	-8	-8	-1	-12	-20	-51	-97	-101	-30	-78	-23	-276	-367	-442	-466	-224	-81

anticipated the direction of future cumulative correlations in real time. Investors have access only to backward-looking information, so the results avoid look-ahead bias by using a pseudo out-of-sample approach. In each period, I form two forecasts, with the null assuming zero cumulative correlations at all lags and the alternative assuming correlations that reflect the history up to that point. The null hypothesizes iid returns while the alternative assumes past lead-lag patterns are persistent and repeat themselves. I compare these two forecasts to actual realized future led-lags using next 12-month returns for monthly results and next 10-year returns for annual results. Squared forecast errors are then calculated, and the exercise is repeated on a one-step-ahead basis for the next month or year. Correlation forecasts are updated recursively at every step using the latest available information given an expanding historical window. When squared forecast errors are calculated for all periods, they are averaged to form the mean-squared-error (MSE) for each forecast. Panel (b) shows the OOS R^2, where positive numbers mean the alternative hypothesis has outperformed while negative ones reflect null out-performance. Out-of-sample predictability is rare compared to in-sample predictability, and Goyal and Welch (2008) show that nearly all market mean-return predictors lack out-of-sample performance. Nevertheless, panel (b) shows that positive OOS R^2 occur at the monthly horizon for market autocorrelations and at the annual horizon for market-SMB correlations at two years. This indicates strong persistence in some lead-lag patterns such that they are possibly actionable in real-time. For SMB especially, investors could have anticipated long-horizon betas that deviate from short-horizon ones and formed adjusted conditional expectations accordingly.

4 The Shape of Alpha and Beta Average Term Structures

This section presents the paper's main empirical findings and shows the average term structure shape for size, value, and momentum. For each portfolio and for alpha and beta separately, I compare the three methods described in the previous section: unconditional, conditional rolling-window, and the new conditional high-frequency approaches. **Table 2** displays coefficient estimates and corresponding Newey-West '94 standard errors. The three approaches are broadly consistent.

Table 2: Alpha and Beta Term Structures. Data from 1926 - 2015 used to form non-overlapping monthly, quarterly, annual returns and annual-overlapping 3-year, 10-year returns. Tables show unconditional and average conditional alpha and beta term structures. Unconditional alphas and betas are calculated using all historical data in an OLS regression: $r^i_{t\to t+h} = \alpha^U_{t\to t+h} + \beta^U_{t\to t+h} r^m_{t\to t+h}$. Conditional alpha and beta methods described in Section 2. T-statistics are shown in brackets using Newey-West '94 standard errors. Boldfaced coefficients denote a significant difference from zero at a two-sided 95% level.

(a) Size (SMB)

	Alpha					Beta				
	Mn	Qt	Yr	3Y	10Y	Mn	Qt	Yr	3Y	10Y
Unconditional	0.9	0.4	0.7	1.2	**3.9**	**0.2**	**0.3**	**0.2**	0.2	-0.2
	[0.9]	[0.4]	[0.6]	[0.5]	[3.1]	[6.0]	[4.6]	[3.6]	[1.3]	[-1.7]
Cond- Rolling	**1.3**	1.3	0.0	1.0	**4.7**	**0.2**	**0.2**	**0.3**	0.2	-0.2
Window	[10.8]	[1.9]	[0.0]	[1.1]	[2.4]	[8.6]	[8.6]	[5.6]	[0.4]	[-0.7]
Cond- High	**1.9**	1.0	0.6	**2.0**	**5.0**	**0.1**	**0.2**	**0.3**	0.0	**-0.3**
Frequency	[7.4]	[1.5]	[0.7]	[2.2]	[6.1]	[4.6]	[6.7]	[4.5]	[0.2]	[-2.8]

(b) Value (HML)

	Alpha					Beta				
	Mn	Qt	Yr	3Y	10Y	Mn	Qt	Yr	3Y	10Y
Unconditional	**3.1**	**3.1**	**3.5**	**3.8**	**4.9**	0.1	0.1	0.1	0.0	-0.1
	[2.2]	[2.2]	[2.8]	[2.5]	[7.9]	[1.8]	[1.1]	[0.8]	[0.4]	[-0.5]
Cond- Rolling	**3.8**	**4.2**	**3.8**	**4.3**	**4.9**	-0.0	-0.0	0.0	-0.0	-0.0
Window	[13.1]	[8.4]	[5.3]	[5.7]	[4.3]	[-0.1]	[-0.3]	[0.1]	[-0.0]	[-0.0]
Cond- High	**4.7**	**4.3**	**3.6**	**4.4**	**5.2**	-0.0	0.0	0.1	-0.0	-0.1
Frequency	[19.3]	[10.0]	[4.3]	[7.9]	[6.9]	[-0.5]	[0.2]	[0.6]	[-0.1]	[-0.4]

(c) Momentum (UMD)

	Alpha					Beta				
	Mn	Qt	Yr	3Y	10Y	Mn	Qt	Yr	3Y	10Y
Unconditional	**8.6**	**9.1**	**7.4**	**6.1**	5.1	**-0.3**	**-0.4**	-0.1	0.1	0.3
	[5.4]	[5.6]	[3.7]	[2.8]	[1.8]	[-2.4]	[-2.6]	[-1.5]	[1.0]	[1.5]
Cond- Rolling	**6.2**	**7.4**	**7.8**	**6.4**	3.5	**-0.1**	**-0.2**	-0.1	0.1	0.3
Window	[14.2]	[8.4]	[5.1]	[3.8]	[0.2]	[-4.4]	[-3.8]	[-1.9]	[1.1]	[1.4]
Cond- High	**6.0**	**7.0**	**7.2**	**7.4**	4.5	**-0.0**	-0.1	-0.1	-0.0	0.2
Frequency	[12.1]	[7.1]	[6.3]	[6.0]	[1.3]	[-2.1]	[-1.0]	[-0.7]	[-0.6]	[1.1]

4.1 Discussion of Size, Value, and Momentum Results

Size's monthly betas are positive and strongly significant. They increase and peak at the quarterly or annual horizon, and then decline until they're negative by the 10-year horizon. Only the high-frequency approach shows statistical significance. Beta term structures impact alpha term structures directly. They are positive at the monthly horizon and significantly so under both conditional approaches. As the horizon lengthens, alpha changes follow beta ones but with the opposite sign, falling as betas rise and then rising as betas fall. Conditional alphas shrink and lose their significance at the quarterly and annual horizons. Handa, Kothari, and Wasley (1989) and Handa, Kothari, and Wasley (1993) find that the size anomaly disappears under annual instead of monthly returns. They examine only unconditional results over a more limited time period, however, and I provide updated, confirming evidence for conditional alphas as well. In addition and more strikingly, the pattern reverses at multi-year horizons, with alphas becoming highly significant by 10-years with magnitudes roughly triple that of monthly ones. To my knowledge, this is the first time this pattern has been documented. Recent work (e.g. Hou and Van Dijk (2014), Asness et al. (2015)) has shown ways to resurrect the size effect given its poor performance in the past three decades. It turns out that long-horizon returns are another place where the size effect is alive and well.

From equation (4), the mean conditional alpha term structure is determined not just by beta effects but also by the covariance between conditional betas and the market risk-premium (i.e. the market's conditional mean). **Figure 4** plots the covariance's term structure on the same scale as **Figure 1**. Under both conditional approaches, the level is near zero and the term structure is essentially flat. This term thus plays no major role in either the level or shape of the alpha term structure, and it must be beta effects that drive alpha results. The result is consistent with Lewellen and Nagel (2006)'s and Ang and Kristensen (2012)'s findings that this covariance term is simply too small to explain the magnitude of these anomalies' alpha.

In the results presented for size so far, I've only demonstrated the significance of coefficients in their deviation from zero. For betas, this is important for understanding whether the portfolios are risks or hedges, and for alphas, whether they are anomalous with respect to the CAPM. However, I have not yet shown that coefficients at different horizons differ from each other in a statistically significant way and therefore, the term structure is indeed sloped. Looking at **Figure 1**, we can get a sense of this by examining whether the 95% confidence interval at different horizons overlap with each other. No overlap suggests significance but this approach is too stringent since it ignores the likely correlation of coefficient estimates across horizons. For example, a high monthly beta estimate for a particular period is likely to occur during a high annual estimate too. This means coefficient

Figure 4: Term Structure of Covariance Term between Conditional Betas and the Market Risk Premium. Data from 1926 - 2015 used to form non-overlapping monthly, quarterly, annual returns and annual-overlapping 3-year, 10-year returns. Conditional alpha and beta methods described in Section 2. Figures show the unconditional covariance between conditional betas and the market risk premium (i.e. the market conditional mean return estimated using the rolling-window approach).

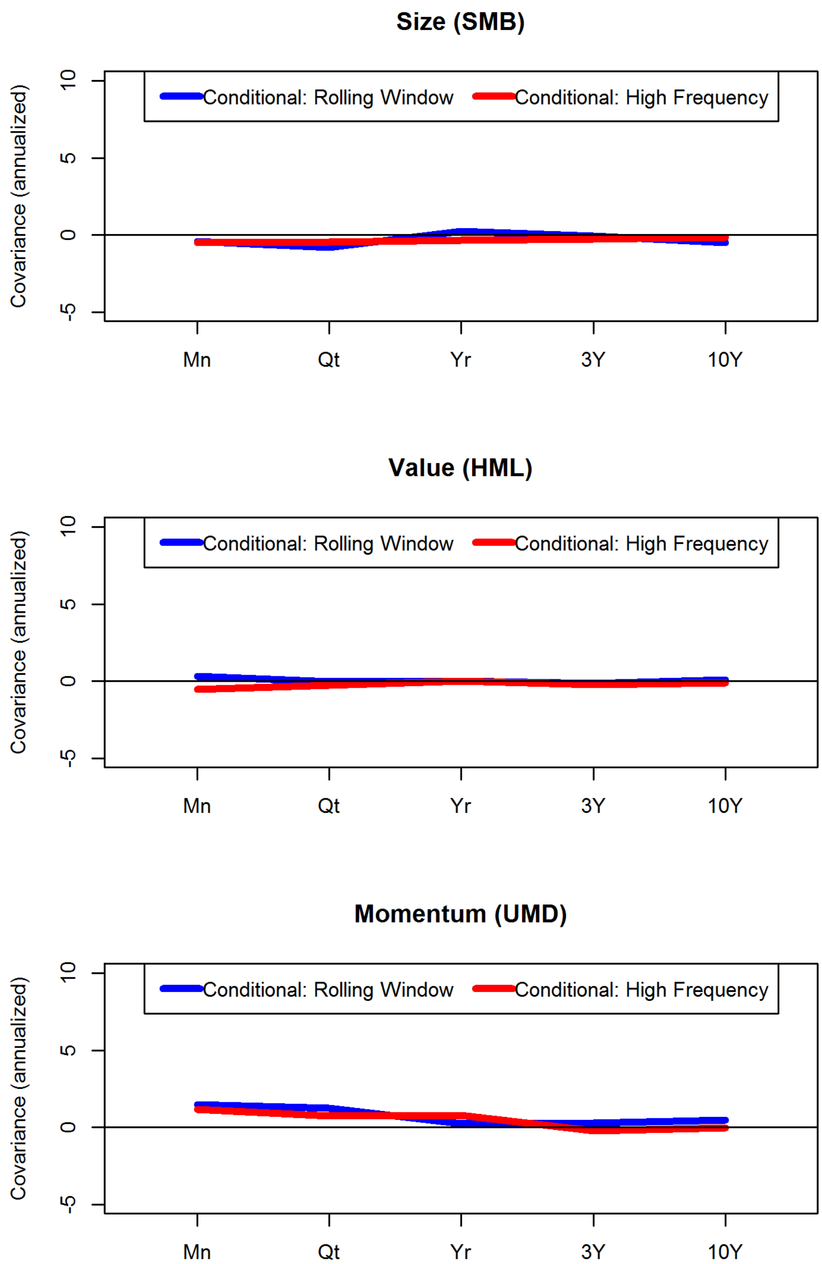

differences across horizons can be small yet still be significant. To account for this possibility, I continue using a t-test to assess significance but incorporate in the denominator standard error correlations. These are estimated with an exactly-identified GMM setup for each pairwise horizon. GMM estimates the coefficient covariance matrix that's easily convertible to correlations and by using a Newey-West '94 kernel, leads to estimates robust to heteroskedasticity and autocorrelation.

Table 3 presents t-statistics of horizon coefficient differences that incorporate GMM correlation estimates. Positive values imply an increase in alphas or betas as the horizon lengthens. Looking only at size for now, I find significant coefficient differences primarily for 10-year versus shorter horizons. Betas are significantly more negative while alphas are significantly more positive. The new high-frequency approach shows the strongest results, especially for betas. Alpha significance is rarer because standard errors include variation from both betas and expected returns. I obtain only significance for alphas increases at multi-year horizons and not for alpha decreases going from monthly to quarterly or annual horizons. Prior works' finding that the size effect becomes less anomalous at annual horizons is thus much weaker than and completely reversed by the opposite effect at multi-year horizons. Size becomes *more* anomalous at these longer horizons, and it is only here are differences significant.

Turning to value, **Table 2** shows consistent results among all three estimation approaches. We see insignificant betas and highly significant alphas across all examined horizons. Value's beta is indistinguishable from zero regardless and given value's excess return historically, translates to significant alphas at all horizons. Alpha's flat term structure reflects beta's flat-term structure since the covariance term between beta and the market-premium plays a negligible role (**Figure 4**). **Table 3** contains almost no significant differences between any two horizons for all estimation approaches. The null hypothesis that the value term structures are flat and that different horizon returns result in the same alpha and beta estimates thus stands. These patterns are unsurprising given earlier findings showing value's lack of significant lead-lag correlations (**Table 1**). Unlike size or momentum, value lacks monthly or annual lead-lag correlations that consistently push in the same direction. Given this necessary condition for sloped average term structures, value's alpha and beta stay similar regardless of the investment horizon.

These results contrast with R. Cohen, Polk, and Vuolteenaho (2009) who find that value stocks

Table 3: T-Statistics for Alpha and Beta Term Structure Differences. Data and method as in **Table 2**. Table shows the t-statistic of alpha and beta differences across different horizons (denoted in the rows versus columns). For example, $Tstat(\beta_{Qt} - \beta_{Mn}) = \frac{\beta_{Qt}-\beta_{Mn}}{\sqrt{se(\beta_{Qt})^2+se(\beta_{Mn})^2-2\rho_{Qt,Mn}se(\beta_{Qt})se(\beta_{Mn})}}$ where $\rho_{Qt,Mn}$ denotes the correlation between the coefficient estimates calculated using pairwise GMM. Positive numbers denote an increase in alpha or beta with a lengthening of the horizon. Boldfaced tstats denote a significant difference at a two-sided 95% level.

(a) Size (SMB)

		Alpha				Beta			
		Qt	Yr	3Y	10Y	Qt	Yr	3Y	10Y
Unconditional	Mn	-0.3	-0.2	0.1	1.7	1.2	0.6	-0.2	**-2.9**
	Qt		0.1	0.3	**2.0**		-0.4	-0.7	**-3.1**
	Yr			0.3	1.7			-0.6	**-3.3**
	3Y				1.0				**-3.3**
Cond- Rolling Window	Mn	-0.0	-2.0	-0.4	1.7	1.3	1.5	0.0	-1.4
	Qt		-1.4	-0.2	1.4		1.0	-0.1	-1.7
	Yr			1.4	**2.0**			-0.1	-1.8
	3Y				1.6				-0.9
Cond- High Frequency	Mn	-1.2	-1.5	0.1	**3.5**	**3.4**	**3.0**	-0.9	**-3.9**
	Qt		-0.4	0.8	**3.5**		0.8	**-2.0**	**-4.8**
	Yr			1.4	**4.0**			**-2.5**	**-4.8**
	3Y				**3.1**				**-2.7**

(b) Value (HML)

		Alpha				Beta			
		Qt	Yr	3Y	10Y	Qt	Yr	3Y	10Y
Unconditional	Mn	-0.0	0.2	0.3	1.1	0.1	-0.5	-0.6	-1.3
	Qt		0.2	0.4	1.2		-0.5	-0.6	-1.3
	Yr			0.4	1.2			-0.3	-1.1
	3Y				0.8				-0.9
Cond- Rolling Window	Mn	0.7	0.1	0.8	1.1	-0.2	0.1	-0.0	-0.0
	Qt		-0.4	0.1	0.5		0.2	0.0	-0.0
	Yr			**2.5**	1.4			-0.1	-0.1
	3Y				0.9				-0.0
Cond- High Frequency	Mn	-0.9	-1.4	-0.5	0.6	0.4	0.9	0.1	-0.3
	Qt		-0.7	0.2	1.3		0.4	-0.2	-0.4
	Yr			1.7	1.4			-1.2	-0.8
	3Y				0.9				-0.3

(c) Momentum (UMD)

		Alpha				Beta			
		Qt	Yr	3Y	10Y	Qt	Yr	3Y	10Y
Unconditional	Mn	0.2	-0.4	-0.9	-1.1	-0.5	1.2	**2.6**	**2.7**
	Qt		-0.6	-1.1	-1.2		1.6	**2.8**	**3.0**
	Yr			-0.7	-0.9			**2.2**	1.9
	3Y				-0.7				1.4
Cond- Rolling Window	Mn	1.2	1.2	0.2	-0.2	-1.4	0.7	**2.8**	1.9
	Qt		0.2	-0.4	-0.2		1.7	**4.3**	**2.6**
	Yr			-1.0	-0.3			1.7	1.7
	3Y				-0.2				1.4
Cond- High Frequency	Mn	0.9	1.0	1.1	-0.4	-0.1	-0.3	0.0	1.3
	Qt		0.1	0.2	-0.6		-0.2	0.1	1.2
	Yr			0.2	-0.9			0.5	1.2
	3Y				-1.1				1.2

have rising betas and declining alphas over multi-year horizons. Our approaches to value investing, however, are different. R. Cohen, Polk, and Vuolteenaho (2003) hold a fixed basket of stocks that start out as value but may drift over time and not be considered value after a few years. In contrast, I hold a value *strategy* that automatically rebalances toward new value stocks each year. R. Cohen, Polk, and Vuolteenaho (2003) calculate cashflow betas whereas my betas reflect period-end total returns from all sources, including dividends assumed to be reinvested along the way. Their results are thus more relevant for holding an initially-cheap asset for a long time whereas mine are more applicable for implementing a consistent value stock investing strategy. Since the portfolios differ, estimation results may differ too. Nevertheless, one central point of agreement is that alpha and beta estimates can vary depending on the investment horizon.

Momentum results tend to be the reverse of size's but weaker. Monthly betas are significantly negative but increase and turn positive at multi-year horizons. Alphas are strongly significant at short and medium term horizons but decline and lose significance by 10-years. Increasing betas drive down alphas but widening standard errors also contribute. Indeed, investing over the long run is a rare example of momentum not being anomalous with respect to the CAPM. **Table 3** shows significant differences between multi-year betas and shorter horizons under the unconditional and conditional rolling-window approaches. Alpha differences, however, exhibit no significance at all. Momentum alphas do decline substantially, with magnitudes comparable to size's alpha changes. But momentum 10-year standard errors are much wider, making it harder to draw significant inferences. Like size and value, **Figure 4** also shows that the term structure of the covariance term between momentum beta and the market risk-premium is mostly flat. There is a slight ~1% elevation at shorter horizons, suggesting some correlation without which short-term alphas would be even higher. Positive correlations increase the market's ability to explain momentum's excess returns and thus lower the anomalous alpha. Since the correlation diminishes at longer horizons, this also dampens the degree to which the alpha term structure slopes downward.

The paper's momentum results are consistent with Daniel and Moskowitz (2015). They find that momentum crashes tend to follow large market declines and coincide with subsequent market rebounds. This pattern is partly forecastable, and proper risk management can lead to much better momentum performance. Negative momentum returns in a rising market environment contribute to negative contemporaneous short-term betas. When this follows market declines, they add to positive lagged correlations. These patterns result in an upward-sloping momentum beta term structure, which is exactly what I find. Longer horizons pick up lead-lag patterns short-horizons avoid, including the tendency for momentum to crash *after* a market crash. This positive correlation thus raises beta and lowers alpha. The momentum strategy I study assumes automatic rebalancing toward momentum regardless of market conditions. Daniel and Moskowitz (2015) as well as Barroso and Santa-Clara (2015) argue that this isn't optimal when crashes may be anticipated. But this paper's purpose is not to develop optimal trading strategies but rather to document the potential for risk-return expectations to vary according to investment horizon. The forecastability of momentum crashes thus implies that investors can form conditional expectations about these lead-lag correlations and do so on an ex-ante basis.

4.2 Consistency and Implication of Results

Overall, the unconditional and two conditional approaches lead to similar conclusions regarding average term structure shapes. Agreement between the two conditional approaches is reassuring since both should be measuring time-varying ex-ante expectations. But what about agreement between the unconditional and conditional approaches? Focusing on beta's numerator, unconditional and average conditional results are linked through the law of total covariance: $\mathrm{Cov}(r_t^i, r_t^m) = \mathrm{E}[\mathrm{Cov}_t(r_t^i, r_t^m)] + \mathrm{Cov}(\mathrm{E}[r_t^i], \mathrm{E}[r_t^i])$. The unconditional covariance of returns must equal the average conditional covariance of returns *plus* the covariance of the returns' conditional means. Differences in the shape of unconditional and average conditional covariance term structures must, therefore, require a sloped term structure for this additional component. The covariance of the conditional means must vary with the investment horizon, and as already seen in equation (2), this requires them to exhibit lead-lag correlations. But conditional means have processes more stable than the returns they are part of, making lead-lag correlations that much more difficult to detect and estimate. Term structure effects in conditional means are thus likely swamped by term structure effects in returns, making minimal the wedge between unconditional and mean conditional covariance results. Of course, the same analysis also applies to beta's denominator. The finding of broadly consistent unconditional and conditional beta term structure shapes is thus unsurprising. This along with minimal term structure effects in the covariance term between beta and the market-premium (**Figure 4**) makes unsurprising consistent unconditional and conditional alpha term structure shapes too.

The two conditional approaches have similar average alpha and beta coefficients, but Newey-West standard errors for these mean estimates can differ. Given t-statistics in **Table 2**, we can infer that the new high-frequency approach has much smaller standard errors at multi-year horizons. By construction, the rolling-window approach has more persistent estimates over time (see Appendix for summary statistics), and this high autocorrelation makes precise estimation of its mean difficult. On the other hand, the high-frequency approach uses only data within each period, producing independent observations whose mean can be more confidently estimated. Take the longest 10-year horizon with the given 89-year dataset. The rolling-window approach has effectively 2 independent 40-year observations (base case window for 10-year horizons) while the high-frequency approach has nearly 9. Long horizons thus limit the rolling-window approach's usefulness since the approach requires sampling across adjacent periods. Thankfully, the new high-frequency approach samples within periods instead, making it an especially useful method for long-horizon alpha and beta estimation.

The average term structure patterns presented here have important implications for investors implementing size, value, or momentum strategies or equivalently, for investors holding stocks weighted toward SMB, HML, or UMD factors. Size should be considered risky from a market exposure perspective but only at horizons less than one-year. At multi-year investment periods, size adds no risk and in fact, can be considered a hedge when its beta turns negative. Momentum shows the opposite pattern while value carries no significant market exposure regardless of the horizon. These beta patterns produce a size alpha that increases with the horizon, a value alpha that stays similar, and a momentum alpha that decreases. By the 10-year horizon, all three alphas have similar magnitudes of around 4-5% annually, with only momentum lacking statistical significance. Given that size is generally considered weak while momentum strong (e.g. Fama and French (2012), Novy-Marx and Velikov (2016), Harvey, Liu, and Zhu (2015)), this is, as far as I know, the first paper to document how size can outperform momentum on a risk-adjusted basis. You only need to hold the strategies long enough.

5 Stationarity of Conditional Alpha and Beta

The paper has so far focused exclusively on the *average* term structure across all time periods and found some to be significantly sloped. This can only be induced by persistent lead-lag correlations, and I've thus only investigated the first driver of term structure effects. In this section, I cover the second driver, namely the stationary nature of conditional alphas and betas. Stationarity implies mean-reversion after stochastic shocks, such that high alphas or betas tend to be followed by lower ones. In the absence of lead-lag correlations, the long-horizon conditional beta approximates the average conditional short-horizon beta over the horizon (Equation (2)). This means abnormally high short-horizon betas at a particular time tend to occur along with downward-sloping conditional term structures. Given the difficulty in estimating conditional moments precisely, however, it's hard to demonstrate that the term structure at a particular time is sloped in a statistically significant way. As seen in the previous section, long-horizon standard errors are already large for *average* conditional alphas and betas, leading to statistically significant *average* differences across horizons that are challenging to achieve. Obtaining significant differences across horizons for a term structure at a particular point in time would thus be even more difficult. The standard errors are simply too large for meaningful inferences.

I overcome this difficulty by focusing on the demonstration of alpha and beta stationarity. Sloped conditional term structures are then a direct implication, thus avoiding the need to demonstrate this explicitly for individual cases at different time periods. I conduct Augmented Dickey-Fuller tests to assess the null that alphas or betas have unit-roots such that rejection of the null implies stationarity. **Table 4** shows the resulting test statistics under both conditional estimation methods and under alphas and betas estimated using monthly, quarterly, and annual returns. The rolling-window approach's base case window length has 60, 40, and 20 observations, respectively. Even if the true conditional moments are stationary, the way in which the approach's estimates are formed makes rejecting the null especially difficult. Therefore, I explore an alternative specification where the window length is half that of the base case.

Table 4: Augmented Dickey-Fuller Statistics to Test for Unit Roots In Conditional Alphas and Betas. Data from 1926 - 2015 used to form non-overlapping monthly, quarterly, and annual returns used to estimate conditional alphas and betas (methods described in Section 2). Rolling Window (Half) uses 2.5, 5, and 10-year windows to estimate monthly, quarterly, and annual moments, respectively (these are half that of the base case specification). The table shows Augmented Dickey-Fuller test statistics using intercept but no trend and with BIC-specified lags. Boldfaced values denote a significant rejection of the unit root null at the 95% level.

		Alpha			Beta		
		Mn	Qt	Yr	Mn	Qt	Yr
Size	Rolling Window	-2.5	**-4.0**	**-4.1**	-2.7	**-4.5**	-0.8
(SMB)	Rolling Window (Half)	**-3.6**	**-4.9**	**-3.0**	**-3.6**	**-4.2**	-2.4
	High Frequency	**-4.3**	**-3.7**	**-4.6**	**-11.7**	**-12.3**	**-6.4**
Value	Rolling Window	**-4.2**	**-3.2**	-2.0	-1.6	-1.6	-0.9
(HML)	Rolling Window (Half)	**-4.9**	**-3.7**	-2.1	-2.3	-1.9	-2.0
	High Frequency	**-6.4**	**-9.2**	**-5.6**	**-6.4**	**-8.9**	**-5.5**
Momentum	Rolling Window	**-3.6**	**-3.4**	-1.8	-2.6	**-3.0**	-1.4
(UMD)	Rolling Window (Half)	**-4.7**	**-4.3**	-2.2	**-3.5**	**-3.5**	-2.4
	High Frequency	**-5.5**	**-5.1**	**-4.2**	**-11.8**	**-9.7**	**-5.0**

Under the new high-frequency method, significant rejection of unit-roots occurs in all cases. Under the rolling-window approach, they occur in most cases, with HML betas and annual moments being the main exceptions. As expected, halving the window length leads to stronger results, but the findings are broadly consistent with base case windows. Consequently, stationary alphas and betas imply sloped conditional term structures when these moments receive an especially large shock or when they are estimated with especially large errors. Even in the absence of lead-lag correlations, the conventional practice of using conditional monthly alphas and betas for long-term investors is thus incorrect. In the absence of these correlations, long-horizon conditional moments are better approximated by *unconditional* rather than *conditional* short-horizon moments. The long-term investor is better served with regression estimates that use all historic data rather than the most recent rolling-window. Using recent data estimates conditional results that lack persistence while using all data better captures the long-horizon after conditional information gradually dissipates. Therefore, even if the evidence presented in the previous section for lead-lag correlations and for sloped *average* term structures is unpersuasive, sloped *conditional* term structures still occur because of stationary conditional alphas and betas.

6 Conclusion

This paper examines the unconditional and conditional term structures of CAPM alphas and betas for long-short portfolios sorted on size, value, and momentum. The literature traditionally uses monthly returns to estimate alphas and betas even though they are inappropriate for long-term investors. Many investors have horizons that span years and are more interested in long-horizon moments. In addition, firms' investing or financing decisions use discount rates matched to cashflows and thus require betas of the right horizon. Researchers should thus explore alternative return periods and not rely on monthly returns exclusively. More generally and especially since the financial crisis, there's broad acknowledgment of the pitfalls associated with Wall Street's short-term thinking. Looking at returns beyond just 1-month is perhaps a small step away from that.

Long-horizon betas can deviate from short-horizon ones for two reasons. First, portfolios exhibit lead-lag correlations that are significant and persistent across time. Long-horizon betas can be dominated by these correlations rather than simply reflect short-term contemporaneous comovement. I find this especially true for size, which exhibits strongly negative annual lead-lag correlations with the market. This leads to size's beta sign reversal such that size turns from a risk into a hedge at multi-year horizons. Value has lead-lag effects that are weak and largely cancel, resulting in a mostly flat beta term structure. Momentum is the opposite of size but with weaker horizon effects. The second reason for sloped beta term structures is that short-horizon conditional betas are stationary and mean-reverting. Abnormally large short-horizon betas tend to be followed by smaller ones such that a period-specific conditional term structure can be sloped even in the absence of lead-lag correlations.

These beta term structure effects directly impact alphas. A higher beta that explains more of the portfolio's excess return will thus diminish unexplained alpha. The other possible alpha driver, the covariance term between conditional beta and the market risk-premium, shows no term structure dynamics. Therefore, size's downward-sloping beta term structure translates into an upward-sloping alpha term structure. Momentum's is the opposite while value's is flat. At the 10-year horizon, all three portfolios have comparable alphas, and it is momentum's that's least significant. Horizon does matter.

Many questions remain unanswered. First, what are the economic drivers behind these significant and persistent lead-lag correlations? If slow information diffusion can explain positive monthly correlations for the size portfolio, what can explain the much larger and *negative* annual correlations? Second, are investors aware of these term structure effects, and do they take advantage of them? For example, do long-term investors capitalize on the upward slope of size's alpha term structure and invest more in size relative to short-term investors? Finally, what are the general equilibrium implications of these empirical findings? How are prices determined given the presence of heterogeneous investors with different horizons? I plan on exploring these issues in future drafts or additional papers.

Appendix I. Summary Descriptions

In this appendix, I provide basic summary statistics with a brief discussion. **Table A1** panel (a)'s left-hand side contains summary statistics for each portfolio's excess return. This is done at horizons ranging from 1-month to 10-years. All returns are log returns, so long-horizon returns simply sum short-horizon ones. Since different horizons use the same historical data, returns of different horizons should all have the same mean statistic when annualized. However, the table shows slight deviations at multi-year horizons because these horizons use overlapping annual data. This effectively underweights the beginning and end portion of the dataset, deviating from the uniform weighting of non-overlapping returns at short-horizons. **Table A1** panel (a)'s right-hand side contains summary statistics for each portfolio's conditional mean return, as estimated using backward-looking rolling windows. Use of backward-looking windows exacerbates mean statistic discrepancies by increasing the differences in how the historical dataset is effectively weighted.

The mean return term structure for log returns should be flat, but the same is not necessarily true for gross returns. Long-horizon gross returns are not the simple sum of short-horizon returns but instead, require geometric compounding. Long-horizon mean returns will thus be smaller than short-horizon ones given volatility and Jensen's inequality. However, I find a nearly flat mean return term structure for gross returns too (results not shown). The Jensen inequality term turns out not to be important because I study only excess-return portfolios. These portfolios have Jensen terms that are driven by long-short leg volatility *differences*, and these differences tend to be small.

Table A1 panel (b) contains summary statistics of the conditional alpha and beta estimates under both the rolling-window and the high-frequency approaches. The rolling window approach smooths across periods whereas the high-frequency approach does not, leading to higher beta autocorrelations of 0.4 ~ 1 versus 0 ~ 0.4. Higher persistence leads to lower volatility. Beta standard deviations (annualized) are 0.2 ~ 0.5 for the rolling-window approach and 0.4 ~ 0.7 for the high-frequency approach. These are sizable time-series fluctuations roughly comparable to those estimated by Lewellen and Nagel (2006). Beta variation reflects not only changing market sensitivities but also changing equity compositions since these excess return portfolios are reconstituted annually (or monthly in momentum's case). As for alphas, the high-frequency approach also has larger standard deviations than the rolling-window approach. Differences here are smaller though since both method's alpha use the same conditional mean estimates.

Figure A1 plots alpha and beta time-series under both conditional approaches. Let's look at the most recent 20-year period to get a sense of their properties. Panel (a) actually begins with market returns and volatility, and I focus on market volatility at the annual horizon. The rolling-window approach responds slowly to real-time events since annual results average over 20-year windows. High-volatility periods during the early 2000s and during 2008-2009 shows up smoothed and with a lag. On the other hand, the high-frequency approach registers two distinct spikes that map directly to the crisis periods, with much larger magnitudes that more appropriately reflect these periods' turmoil. Turning to size, value, and momentum in panels (b)-(d), the high-frequency approach also estimates more dramatic peaks and troughs. During the late 90's boom, for example, momentum's annual beta experiences sharp increases while value's beta plummets. This period saw high-growth technology stocks continuously climb upward during an especially strong bull market. In contrast, the rolling-window approach depicts smoothed estimates less sensitive to return moments' real-time movement. This comparison highlights the advantage of using the new high-frequency approach for alpha and beta estimation.

Table A1: Summary Statistics. Monthly portfolio excess returns from Ken French's website for the period 1927-2015. I convert to log returns (%/year) so long horizon returns are the simple sum of monthly returns. Panel (a) shows summary statistics for portfolio returns and for estimated conditional means using backward-looking windows of 5, 10, 20, 30, and 40 years for monthly, quarterly, annual, 3-year, and 10-year returns, respectively. I permit partial windows in the early history whenever there are at least 2 observations. Monthly, quarterly, and annual returns are non-overlapping while 3-year and 10-year ones are annual-overlapping (autocorrelation statistics use non-overlapping returns in all cases though). Mean statistics are in percent and annualized (e.g. x12 for monthly). Standard deviations are also in percent and annualized but done differently for returns (scale by square root, e.g. x$\sqrt{12}$) versus for conditional means (scale directly, e.g. x12). This accounts for the different time-series properties of returns (little autocorrelation) versus of conditional means (highly persistent). Panel (b) shows summary statistics for conditional alphas (%/year) and betas estimated using two different methods. The rolling-window approach is the OLS regression alpha and beta using a backward-looking window of returns of the same horizon (with same window lengths as conditional mean estimates above). The high-frequency approach uses high-frequency data within each horizon and weights lead-lag covariance terms according to its lead-lag. See Section (2) for details of the conditional estimation approaches. The Rolling-HF Corr row shows the estimate's correlation across the two conditional approaches. All single-lag autocorrelations (Auto1) are adjusted for finite sample bias where ρ is back-solved given $\hat{\rho}$ and the bias term, $E[\hat{\rho} - \rho] = -\frac{1+4\rho}{T}$ (Kendall (1954)).

(a) Portfolio Returns and Conditional Mean Estimates

		Return					Conditional Mean				
		Mn	Qt	Yr	3Y	10Y	Mn	Qt	Yr	3Y	10Y
Obs	No Overlap	1068	356	89	30	9	1067	355	88	29	9
	Yr Overlap				87	80				86	79
Market	Mean	6.0	6.0	6.0	5.7	6.2	6.2	6.0	5.6	4.3	5.9
(MktRF)	SD	18.7	21.2	20.0	19.8	16.0	8.7	6.6	5.3	5.3	2.5
	Auto1	0.1	-0.0	0.1	-0.3	0.1	1.0	1.0	0.8	0.7	0.9
Size	Mean	2.0	2.0	2.0	2.1	2.4	2.0	1.8	1.7	1.8	3.0
(SMB)	SD	10.9	11.8	12.3	14.7	11.2	6.2	4.1	3.2	3.2	1.6
	Auto1	0.1	-0.0	0.3	0.0	-0.5	1.0	1.0	0.9	0.2	1.0
Value	Mean	3.9	3.9	3.9	4.1	4.5	4.1	4.0	3.9	4.1	4.8
(HML)	SD	11.9	13.8	12.3	11.2	7.6	4.9	2.9	2.3	1.7	0.7
	Auto1	0.2	-0.0	-0.0	-0.1	-0.2	0.9	0.9	0.8	0.5	0.4
Momentum	Mean	6.7	6.7	6.7	6.5	6.7	6.9	7.4	7.6	7.0	5.8
(UMD)	SD	18.2	20.6	18.5	17.6	17.0	7.5	6.1	4.8	4.6	4.3
	Auto1	0.1	-0.1	-0.0	-0.2	0.3	1.0	1.0	0.9	0.8	1.0

(b) Conditional Alphas and Betas

			Alpha					Beta				
			Mn	Qt	Yr	3Y	10Y	Mn	Qt	Yr	3Y	10Y
Obs	No Overlap		1067	355	88	29	9	1067	355	88	29	9
	Yr Overlap					86	79				86	79
Size (SMB)	Rolling Window	Mean	1.3	1.3	0.0	1.0	4.7	0.2	0.2	0.3	0.2	-0.2
		SD	1.7	3.3	5.1	5.9	5.5	0.2	0.4	0.2	0.3	0.3
		Auto1	1.0	0.6	0.7	0.1	0.7	1.0	0.4	0.8	0.9	0.8
	High Frequency	Mean	1.9	1.0	0.6	2.0	5.0	0.1	0.2	0.3	0.0	-0.3
		SD	2.3	2.9	5.5	8.9	10.9	0.5	0.5	0.5	0.5	0.5
		Auto1	0.5	0.4	0.3	0.1	0.4	0.1	0.0	-0.0	0.3	0.3
	Rolling-HF	Corr	0.7	0.4	0.4	0.6	0.4	0.3	0.1	0.0	-0.0	0.3
Value (HML)	Rolling Window	Mean	3.8	4.2	3.8	4.3	4.9	-0.0	-0.0	0.0	-0.0	-0.0
		SD	1.5	1.8	2.5	3.3	4.0	0.3	0.3	0.3	0.3	0.3
		Auto1	0.9	0.7	0.8	0.8	1.0	1.0	0.9	1.0	1.0	1.0
	High Frequency	Mean	4.7	4.3	3.6	4.4	5.2	-0.0	0.0	0.1	-0.0	-0.1
		SD	2.0	2.8	5.9	6.1	7.4	0.5	0.6	0.7	0.6	0.4
		Auto1	0.4	0.2	0.3	0.3	-0.3	0.3	0.2	0.3	0.2	-0.0
	Rolling-HF	Corr	0.7	0.4	0.4	0.4	0.3	0.5	0.4	0.4	0.3	0.5
Momentum (UMD)	Rolling Window	Mean	6.2	7.4	7.8	6.4	3.5	-0.1	-0.2	-0.1	0.1	0.3
		SD	2.2	4.3	5.0	9.1	17.8	0.3	0.5	0.2	0.2	0.3
		Auto1	1.0	0.4	0.8	0.6	1.0	1.0	0.4	0.7	0.7	0.4
	High Frequency	Mean	6.0	7.0	7.2	7.4	4.5	-0.0	-0.1	-0.1	-0.0	0.2
		SD	2.9	3.9	7.4	10.8	23.7	0.7	0.7	0.7	0.6	0.7
		Auto1	0.6	0.6	0.3	0.4	0.9	0.4	0.3	0.2	-0.0	-0.0
	Rolling-HF	Corr	0.7	0.5	0.7	0.7	0.6	0.1	0.1	0.4	-0.2	0.1

Figure A1: Times Series of Market Volatility and of Portfolio Alphas and Betas. Data from 1927-2015 used to form non-overlapping monthly, annual returns and annual-overlapping 10-year returns. Panel (a) shows the time-series of market excess-returns (%/year) and of its standard deviation (%/year). Monthly, annual, and 10-year rolling means average the previous 5, 20, and 40-year windows, respectively. Panels (b)-(d) show the time-series plots of size, value, and momentum alphas (%/year) and betas, estimated under the rolling-window and the high-frequency approaches. These methods are described in Section 2.

(a) Market (MktRF)

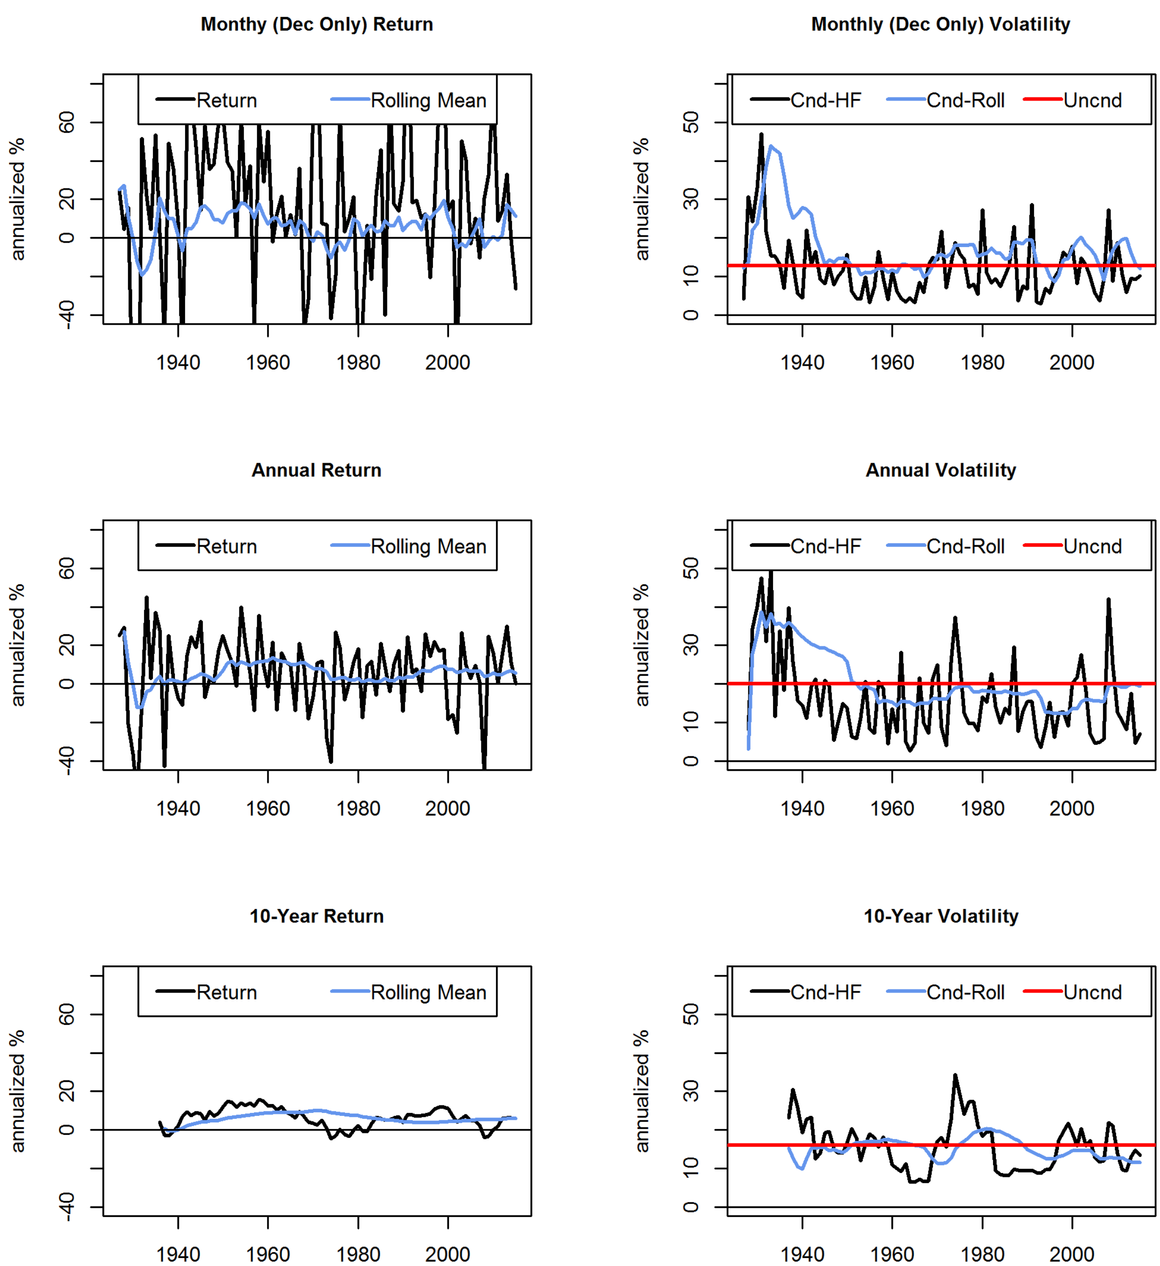

(b) Size (SMB)

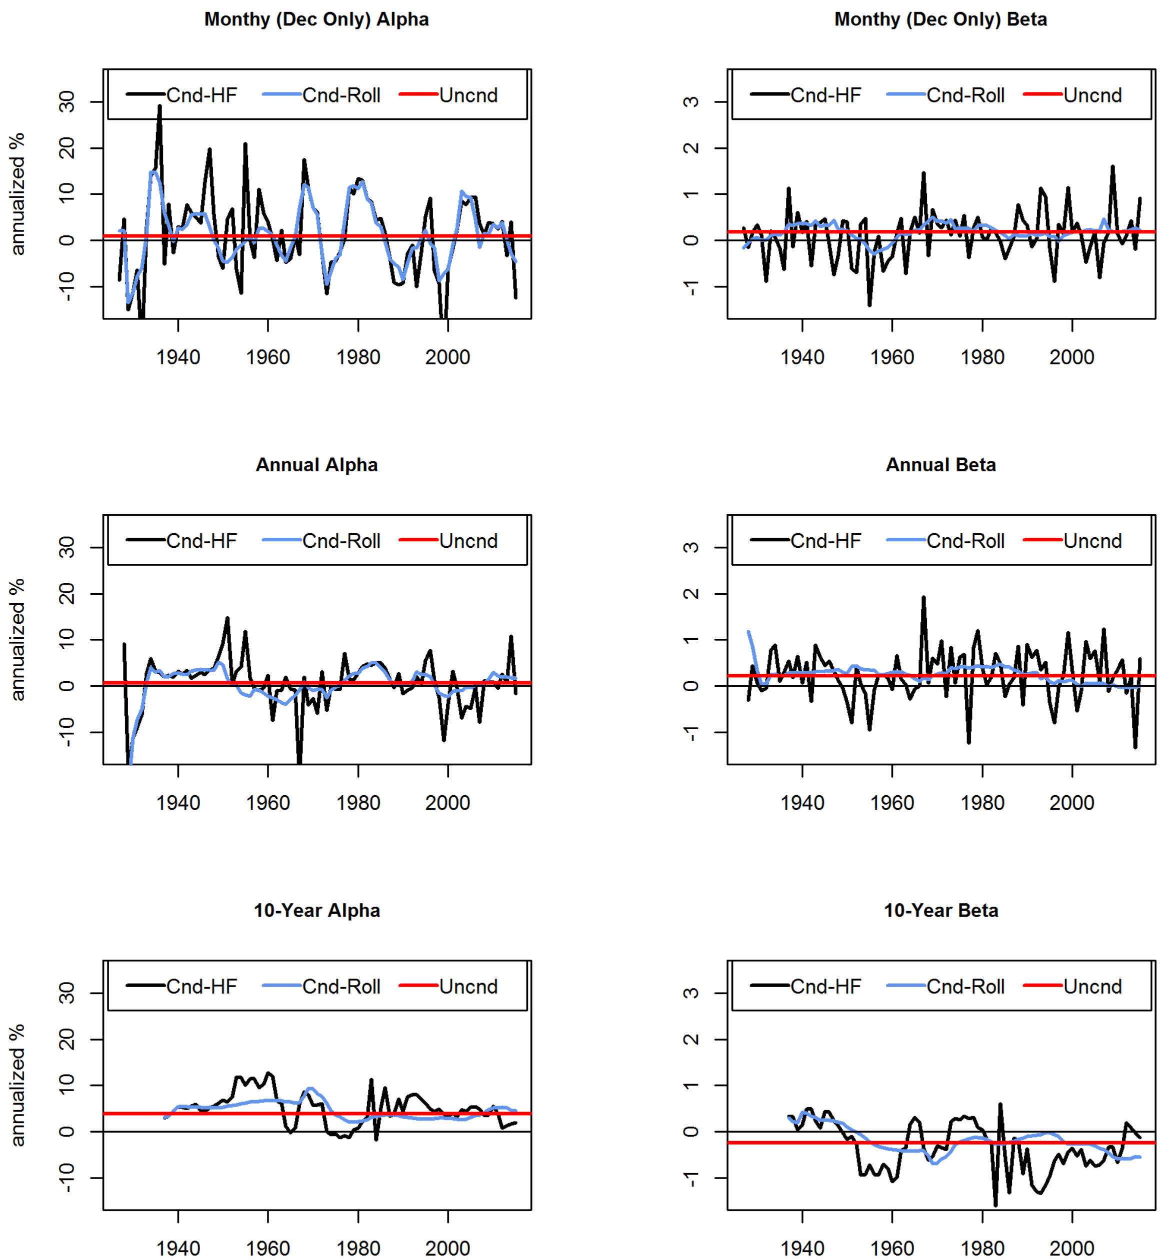
Monthy (Dec Only) Alpha
Monthy (Dec Only) Beta
Annual Alpha
Annual Beta
10-Year Alpha
10-Year Beta
Cnd-HF
Cnd-Roll
Uncnd
annualized %
-10
0
10
20
30
-1
1
2
3
1940
1960
1980
2000

(c) Value (HML)

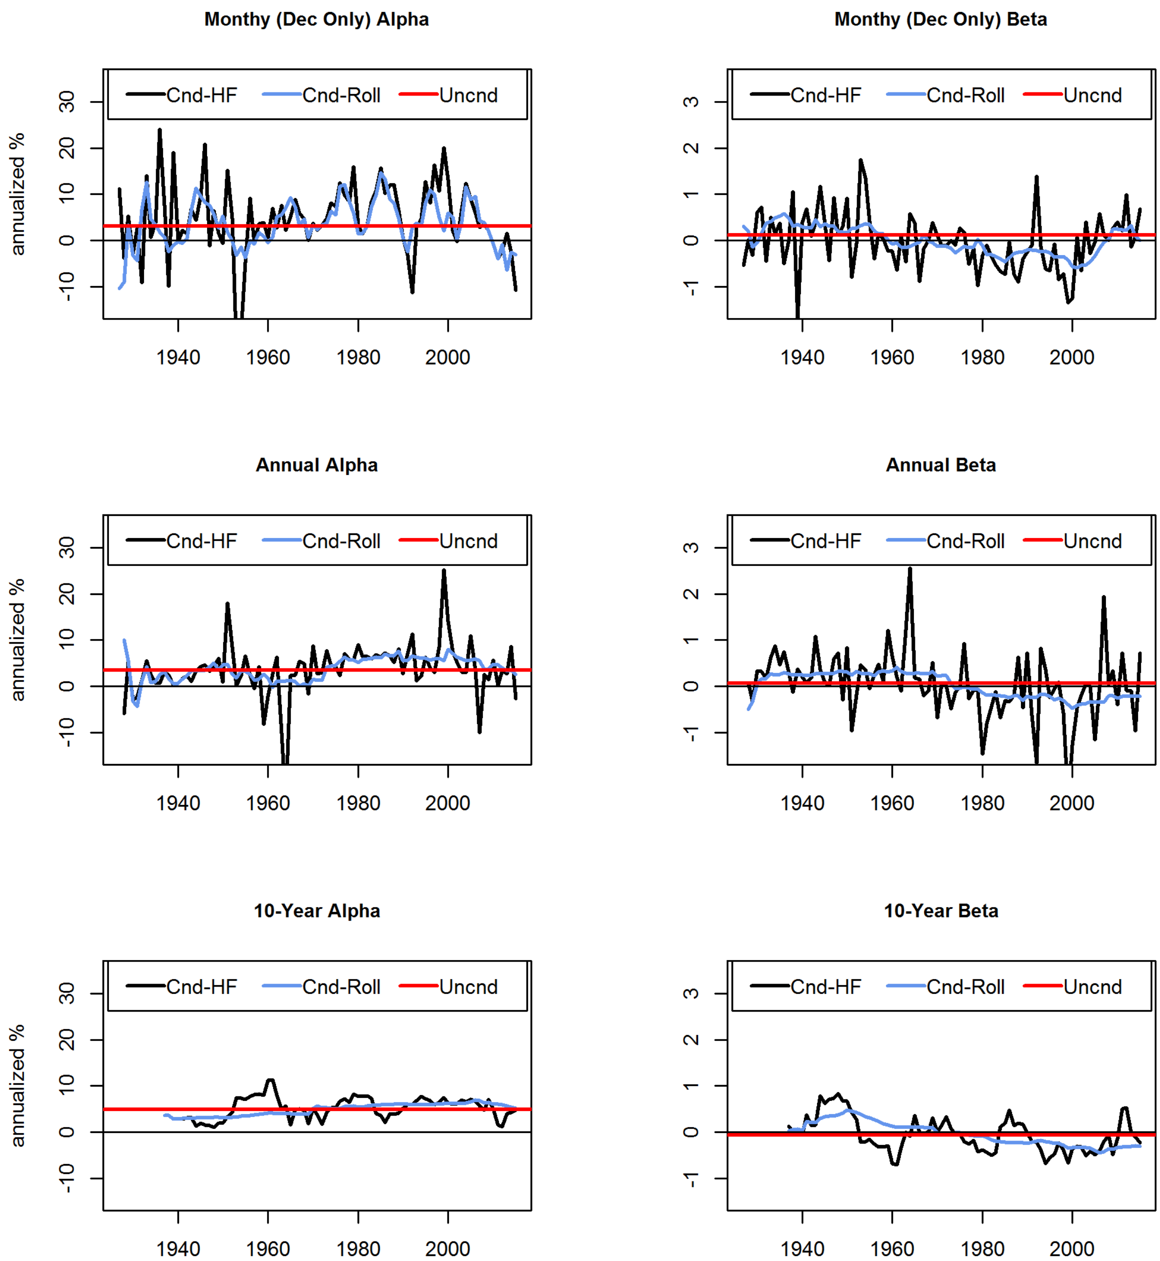
Monthy (Dec Only) Alpha
Monthy (Dec Only) Beta
Annual Alpha
Annual Beta
10-Year Alpha
10-Year Beta
annualized %
Cnd-HF
Cnd-Roll
Uncnd
1940
1960
1980
2000

(d) Momentum (UMD)

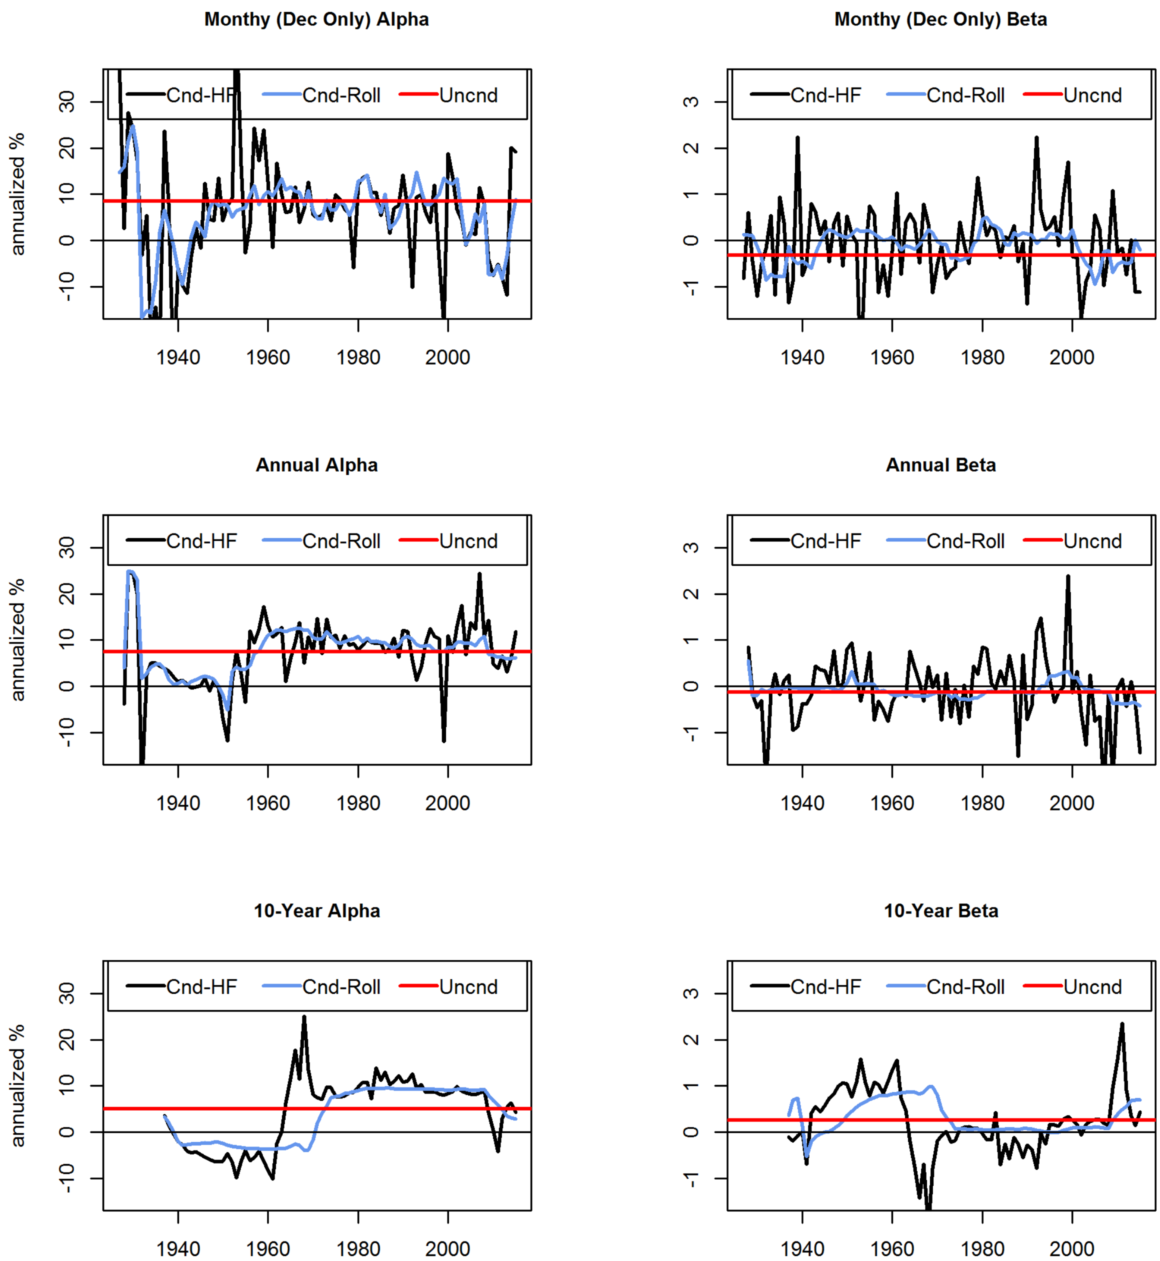

Appendix II. Robustness of Average Term Structure Results

In this appendix, I evaluate the robustness of the main empirical results on the term structure of alphas and betas (**Table 2**). Beginning with the unconditional approach, **Table A2** Panels (a)-(c) presents a series of alternative estimation specifications. I forgo using overlapping annual returns and use only the 30 3-year and 9 10-year observations that are available. I use gross instead of log returns to calculate alphas and betas. To obtain longer-horizon returns, I use annual rather than monthly returns, thus reducing the implicit rebalancing that occurs between the long and short legs. Each alternative involves a single change while preserving all other specifications under the base case. For nearly all coefficients, the sign and degree of significance remain the same, and no alternative specification changes the main findings. In fact, some results strengthen. Size's 10-year beta increases in magnitude and in significance under all alternatives.

Table A2 panel (d) separately examines the long and short legs of the size, value, and momentum excess-return portfolios. This allows identification of the leg driving the results while providing reassurance term structure findings are not artificially induced by the need to short. Size and value are formed annually using the same 6 portfolios, with 3 cuts on book-to-market using NYSE 30th-70th percentiles and 2 cuts on market-value using the NYSE 50th percentile. Size longs an equal-weight of the 3 small-cap portfolios and shorts an equal-weight of the 3 big-cap portfolios. Value longs the 2 value portfolios and shorts the 2 growth portfolios. Momentum is formed monthly, with 3 cuts on 2-12 month prior-returns and 2 cuts on market-value using the same NYSE percentiles. It longs the 2 winner portfolios and shorts the 2 loser ones. See Ken French's website for additional details.

Let's begin with size-value long-only portfolios. Big-Growth stocks are the only portfolio with betas that stay roughly the same across all horizons. Big-Value sees modest declines while Small-Value and Small-Growth see major beta drops of 0.5 and 0.35, respectively. These portfolios have alphas that increase with the horizon, with Small-Value posting the most dramatic gains. The tendency for a downward-sloping size beta term structure and an upward-sloping alpha one is therefore driven by small stocks and to a much lesser extent, value stocks. This is why size's horizon effects are most apparent while value's is slight and insignificant. Let's turn to momentum-size long-only portfolios. The two size portfolios also have the largest horizon effects. These portfolios are on opposing long-short legs, however, so their effects somewhat cancel. Nevertheless, the interaction between size and loser stocks is especially strong. Since momentum shorts this portfolio, its term structures experience the opposite pattern as size's, with rising betas and falling alphas. In sum, small-caps, followed to a lesser extent by loser and value stocks, have the strongest term structure effects, with falling betas and rising alphas as the horizon lengthens. This opens the possibility that a single explanation may account for all the paper's average term structure effects, although the explanation should also clarify why Size-Value and Size-Down interactions are especially strong.

Turning to the robustness of the conditional rolling-window approach, **Table A3** shows results under alternative specifications. In addition to the alternatives explored for unconditional results, different ways of forming rolling-windows are also assessed. First, instead of backward-looking windows, I consider forward-looking ones. Although not implementable in real-time, they may better proxy for investors' conditional expectations under rational expectations. I also consider window sizes that are half of the base case, that use a fixed number of 10 independent observations for all horizons (i.e. 10 months, 10 quarters, etc.), and that use a fixed time period of 20-years for all horizons. Under the base case and all alternatives, I allow for partial windows so that more data can be used. This mainly affects the 10-year horizon and means its 10-Period specification becomes an expanding historical window given its lack of 10e independent observations.

Overall, base case rolling-window patterns are robust to alternative specifications. In nearly all cases, coefficients retain both similar magnitudes and statistical significance thus leading to comparable alpha and beta term structure shapes. I briefly discuss some exceptions. For size betas, the gross return specification does not lead to sign reversal. Nevertheless, there's still a downward-sloping term structure, with an alpha that doubles in magnitude from a monthly to a 10-year horizon. In results not shown, gross returns with non-overlapping data preserve the sign-reversal pattern. For value, look-ahead windows have some alphas that lose statistical significance due to ~1% smaller coefficients. In results not shown, half of this is from higher market-beta covariance and half from lower mean excess returns. The look-ahead approach effectively weights observations at the end of the dataset more, and these have been weaker for value. Finally, momentum's non-overlapping 10-year results are the most anomalous compared to base case ones. The mean beta estimate has as its first observation a partial window that contains only two periods. This leads to an extreme result that swamps others when averaged together. Requiring three instead of just two periods restores the base case pattern. Nevertheless, it's important to acknowledge that 10-year estimates are the least robust for all three portfolios. This is expected and already shows up under the base case, where all three portfolio betas and one alpha lacks statistical significance.

Table A4 completes the robustness section and contains results for the new conditional high-frequency approach. Alternative specifications change the weighting kernel for lead-lag high-frequency returns, using triangular (Bartlett) instead of Gaussian weights or using uniform weighting that includes only the closest half of all lead-lag terms. Other specifications vary the conditional mean assumption, by halving the rolling-window period or by using forward-looking instead of backward-looking windows. I also alter the choice of high-frequency returns, using weekly instead of daily returns for the quarterly horizon, quarterly instead of monthly returns for the 3-year horizon, and annual instead of monthly returns for the 10-year horizon. Finally, I alter the limit in which market autocorrelations can impact the beta denominator, using a denominator minimum of 0.5 or 0.1 versus the base case's 0.3. The main term structure patterns remain robust. The largest changes occur under the alternative of using forward-looking windows to estimate conditional mean returns. Especially for size and for long-horizons, alphas decline substantially, although the significance level and the overall term structure shape remain comparable.

In sum, robustness results add confidence to the term structure patterns shown in **Table 2**. Most specifications are similar to or strengthen base case estimates. Nevertheless, the largest horizon effects do occur at the longest 10-year horizons where uncertainty is indeed the greatest. A skeptical look may discount 10-year results and argue that findings here reflect noise that's obscuring an otherwise flat term structure. I offer four responses to this skepticism. First, as **Table 3** shows, some significant term structure differences occur not just at 10-years but at shorter horizons where data is more abundant. Second, 10-year results are driven by higher-frequency lead-lag correlations with more precise estimates and stronger significance (**Table 1** and **Figure 3**). Lead-lag correlations imply longer-horizon term structure effects even when these effects are statistically insignificant due to a lack of power. Third, this paper builds on a much older and larger literature that has already documented some lead-lag patterns and intervalling effects among stocks. Finally, even in the absence of lead-lag correlations that induce *average* term structure effects, period-specific term structures can still be sloped due to mean-reversion in conditional alphas and betas.

Table A2: Robustness of Unconditional Approach. Data from 1927-2015 used to form non-overlapping monthly, quarterly, annual returns and annual-overlapping 3-year, 10-year returns. The table shows different specifications for estimating unconditional alphas and betas, where estimates are from regressing portfolio excess-returns on the market excess-return using all historical data. *No Overlap* uses non-overlapping returns (affects only 3-year and 10-year horizons which use annual overlapping horizons under the base case). *Gross Returns* uses gross rather than log returns. *Annual Returns* compound annual rather than monthly returns. Panel (d) shows select portfolios for the long and short legs of size, value, and momentum (with risk-free rates subtracted). T-statistics are in brackets using Newey-West '94 standard errors. Boldfaced coefficients denote a significant difference from zero at a two-sided 95% level.

(a) Size (SMB)

	Alpha					Beta				
	Mn	Qt	Yr	3Y	10Y	Mn	Qt	Yr	3Y	10Y
Base Case	0.9	0.4	0.7	1.2	**3.9**	**0.2**	**0.3**	**0.2**	0.2	-0.2
	[0.9]	[0.4]	[0.6]	[0.5]	[3.1]	[6.0]	[4.6]	[3.6]	[1.3]	[-1.7]
No Overlap				1.2	**5.5**				0.1	**-0.6**
				[0.8]	[4.6]				[1.0]	[-3.6]
Gross Returns	1.1	0.3	0.9	2.0	**5.4**	**0.2**	**0.3**	**0.2**	0.2	**-0.2**
	[1.1]	[0.2]	[0.6]	[0.8]	[2.8]	[6.0]	[6.7]	[3.7]	[1.0]	[-2.3]
Annual Returns			0.9	1.6	**4.5**			**0.2**	0.2	-0.3
			[0.7]	[0.6]	[3.5]			[3.9]	[1.1]	[-1.8]

(b) Value (HML)

	Alpha					Beta				
	Mn	Qt	Yr	3Y	10Y	Mn	Qt	Yr	3Y	10Y
Base Case	**3.1**	**3.1**	**3.5**	**3.8**	**4.9**	0.1	0.1	0.1	0.0	-0.1
	[2.2]	[2.2]	[2.8]	[2.5]	[7.9]	[1.8]	[1.1]	[0.8]	[0.4]	[-0.5]
No Overlap				**3.3**	**5.4**				0.1	-0.2
				[3.6]	[5.6]				[0.9]	[-1.3]
Gross Returns	**3.4**	**3.2**	**4.2**	**4.8**	**6.4**	0.1	0.2	0.1	0.0	-0.0
	[2.4]	[2.1]	[3.1]	[2.6]	[6.3]	[1.7]	[1.4]	[0.8]	[0.3]	[-0.2]
Annual Returns			**3.4**	**3.6**	**4.3**			0.1	0.1	0.0
			[2.4]	[2.1]	[4.7]			[0.8]	[0.7]	[0.3]

(c) Momentum (UMD)

	Alpha					Beta				
	Mn	Qt	Yr	3Y	10Y	Mn	Qt	Yr	3Y	10Y
Base Case	**8.6**	**9.1**	**7.4**	**6.1**	5.1	**-0.3**	**-0.4**	-0.1	0.1	0.3
	[5.4]	[5.6]	[3.7]	[2.8]	[1.8]	[-2.4]	[-2.6]	[-1.5]	[1.0]	[1.5]
No Overlap				**4.8**	4.2				**0.3**	0.4
				[2.1]	[1.2]				[2.2]	[1.5]
Gross	**10.5**	**11.3**	**9.6**	**8.6**	**9.4**	**-0.3**	**-0.3**	-0.1	0.0	0.2
Returns	[7.3]	[8.0]	[5.2]	[3.4]	[2.2]	[-2.7]	[-3.6]	[-1.8]	[0.2]	[1.5]
Annual			**8.1**	**6.7**	5.8			-0.1	0.1	0.3
Returns			[4.4]	[2.0]	[0.7]			[-1.1]	[0.5]	[1.0]

(d) Long and Short Legs

	Alpha					Beta				
	Mn	Qt	Yr	3Y	10Y	Mn	Qt	Yr	3Y	10Y
Small Value	2.7	2.4	2.9	3.4	**6.7**	**1.3**	**1.3**	**1.2**	**1.2**	**0.8**
	[1.7]	[1.4]	[1.3]	[1.0]	[3.6]	[21.0]	[16.2]	[12.7]	[6.0]	[3.3]
Small Growth	-2.5	**-3.0**	**-2.9**	-2.4	-0.3	**1.2**	**1.3**	**1.3**	**1.2**	**0.9**
	[-1.8]	[-2.3]	[-3.5]	[-1.1]	[-0.2]	[28.2]	[20.9]	[34.6]	[10.5]	[7.4]
Big Value	0.9	0.8	1.1	1.5	2.4	**1.2**	**1.2**	**1.1**	**1.1**	**1.0**
	[0.7]	[0.7]	[0.9]	[1.1]	[1.7]	[17.9]	[11.9]	[13.6]	[9.0]	[6.1]
Big Growth	-0.2	-0.1	-0.0	-0.1	-0.6	**1.0**	**1.0**	**1.0**	**0.9**	**1.0**
	[-0.3]	[-0.2]	[-0.1]	[-0.2]	[-0.9]	[59.8]	[37.9]	[30.2]	[20.3]	[14.1]
Small Up	**5.7**	**5.5**	**5.6**	**6.0**	**8.3**	**1.2**	**1.2**	**1.2**	**1.2**	**0.9**
	[4.3]	[3.5]	[2.9]	[2.2]	[5.4]	[25.4]	[20.2]	[14.8]	[7.6]	[8.2]
Small Down	**-6.9**	**-7.5**	**-7.2**	-6.2	-3.1	**1.4**	**1.5**	**1.5**	**1.4**	**0.9**
	[-4.7]	[-4.2]	[-3.4]	[-1.8]	[-1.3]	[30.4]	[26.9]	[22.0]	[7.7]	[3.7]
Big Up	**3.5**	**3.7**	**3.3**	**3.2**	**3.1**	**1.0**	**0.9**	**1.0**	**1.0**	**1.0**
	[5.2]	[5.7]	[5.2]	[3.6]	[5.3]	[31.0]	[26.0]	[23.8]	[18.2]	[20.9]
Big Down	**-6.3**	**-6.5**	**-6.1**	**-6.2**	**-4.8**	**1.2**	**1.3**	**1.2**	**1.3**	**1.1**
	[-5.6]	[-5.6]	[-4.8]	[-3.4]	[-5.0]	[17.3]	[16.2]	[14.0]	[8.8]	[8.5]

Table A3: Robustness of Conditional Rolling Window Approach. Data from 1927-2015 used to form non-overlapping monthly, quarterly, annual returns and annual-overlapping 3-year, 10-year returns. The table shows different specifications for estimating the conditional rolling-window approach, where estimates are from regressing portfolio excess-returns on the market excess-return using a backward-looking window. *Base Case* uses windows of 5 years for the monthly horizon, 10 years for the quarterly horizon, 20 years for the annual horizon, 30 years for the 3-year horizon, and 40 years for the 10-year horizon. *No Overlap* uses non-overlapping returns (affects only 3-year and 10-year horizons which use annual overlapping horizons under the base case). *Gross Returns* uses gross rather than log returns. *Annual Returns* compounds annual rather than monthly returns. *Look Ahead* uses a forward-looking window instead of the base case's backward looking window. *Half Window* uses half the window length compared to the base case (i.e. 2.5, 5, 10, 15, 20 years, respectively). *10 Periods* uses 10 non-overlapping observations for each horizon (10 months, 10 quarters, etc.) with partial windows permitted. 20 Years uses rolling windows of 20 years for all horizons. T-statistics are shown in brackets using Newey-West '94 standard errors. Boldfaced coefficients denote a significant difference from zero at a two-sided 95% level.

(a) Size (SMB)

	Alpha					Beta				
	Mn	Qt	Yr	3Y	10Y	Mn	Qt	Yr	3Y	10Y
Base Case	**1.3**	1.3	0.0	1.0	**4.7**	**0.2**	**0.2**	**0.3**	0.2	-0.2
	[10.8]	[1.9]	[0.0]	[1.1]	[2.4]	[8.6]	[8.6]	[5.6]	[0.4]	[-0.7]
No Overlap				1.4	**5.5**				0.1	**-0.4**
				[1.7]	[7.0]				[0.2]	[-3.1]
Gross	**1.6**	**1.3**	0.4	2.5	**3.4**	**0.2**	**0.2**	**0.3**	0.2	0.1
Returns	[12.9]	[2.0]	[0.5]	[1.2]	[3.0]	[8.4]	[8.9]	[5.1]	[0.7]	[1.3]
Annual			0.3	1.6	**5.5**			**0.3**	0.2	-0.2
Returns			[0.4]	[1.6]	[2.1]			[6.5]	[0.3]	[-0.7]
Look Ahead	**1.4**	0.4	0.3	1.6	**4.1**	**0.2**	**0.3**	**0.2**	0.1	**-0.4**
	[6.8]	[0.0]	[0.2]	[0.9]	[10.7]	[13.4]	[4.2]	[2.9]	[0.2]	[-7.1]
Half	**1.1**	1.4	-0.1	1.5	**4.3**	**0.2**	**0.2**	**0.3**	0.1	-0.1
Window	[5.3]	[1.4]	[-0.0]	[0.9]	[2.5]	[11.6]	[7.4]	[3.4]	[0.4]	[-0.4]
10 Periods	**1.3**	1.2	-0.1	1.0	**4.9**	**0.2**	**0.2**	**0.3**	0.2	-0.2
	[4.2]	[1.0]	[-0.0]	[1.1]	[4.1]	[14.9]	[7.3]	[3.4]	[0.4]	[-1.2]
20 Years	**0.8**	**1.0**	0.0	1.4	**4.3**	**0.2**	**0.2**	**0.3**	0.2	-0.1
	[3.4]	[2.2]	[0.0]	[1.1]	[2.5]	[20.5]	[11.2]	[5.6]	[0.4]	[-0.4]

(b) Value (HML)

	Alpha					Beta				
	Mn	Qt	Yr	3Y	10Y	Mn	Qt	Yr	3Y	10Y
Base Case	**3.8**	**4.2**	**3.8**	**4.3**	**4.9**	-0.0	-0.0	0.0	-0.0	-0.0
	[13.1]	[8.4]	[5.3]	[5.7]	[4.3]	[-0.1]	[-0.3]	[0.1]	[-0.0]	[-0.0]
No Overlap				**3.7**	**5.4**				0.0	-0.1
				[3.9]	[5.5]				[0.2]	[-0.5]
Gross Returns	**4.3**	**4.8**	**4.5**	**5.0**	**6.2**	-0.0	-0.0	0.0	-0.1	-0.1
	[13.2]	[8.3]	[4.3]	[3.5]	[5.8]	[-0.0]	[-0.2]	[0.0]	[-0.5]	[-0.7]
Annual Returns			**3.5**	**3.6**	3.8			0.0	0.0	0.1
			[3.3]	[2.5]	[1.1]			[0.1]	[0.0]	[0.1]
Look Ahead	**3.5**	3.2	2.9	**3.4**	4.0	-0.0	-0.0	0.0	-0.1	-0.1
	[11.4]	[0.6]	[1.1]	[4.1]	[0.8]	[-0.3]	[-0.1]	[0.0]	[-0.4]	[-0.1]
Half Window	**3.3**	**3.9**	**3.7**	**4.3**	**5.4**	-0.0	-0.0	0.0	-0.1	-0.1
	[11.0]	[4.3]	[3.4]	[4.0]	[5.5]	[-0.0]	[-0.3]	[0.0]	[-0.4]	[-0.7]
10 Periods	**2.2**	3.1	**3.7**	**4.3**	**4.4**	-0.0	-0.0	0.0	-0.0	0.1
	[5.5]	[1.9]	[3.4]	[5.7]	[3.0]	[-0.0]	[-0.1]	[0.0]	[-0.0]	[0.3]
20 Years	**3.9**	**4.2**	**3.8**	**4.3**	**5.4**	0.0	-0.0	0.0	-0.0	-0.1
	[17.8]	[10.7]	[5.3]	[5.0]	[5.5]	[0.5]	[-0.0]	[0.1]	[-0.2]	[-0.7]

(c) Momentum (UMD)

	Alpha					Beta				
	Mn	Qt	Yr	3Y	10Y	Mn	Qt	Yr	3Y	10Y
Base Case	**6.2**	**7.4**	**7.8**	**6.4**	3.5	**-0.1**	**-0.2**	-0.1	0.1	0.3
	[14.2]	[8.4]	[5.1]	[3.8]	[0.2]	[-4.4]	[-3.8]	[-1.9]	[1.1]	[1.4]
No Overlap				**5.8**	**7.9**				0.2	-0.4
				[2.1]	[4.8]				[1.2]	[-1.1]
Gross Returns	**8.2**	**10.1**	**10.1**	**9.3**	**11.8**	**-0.1**	**-0.2**	**-0.1**	0.0	-0.0
	[31.2]	[13.8]	[6.8]	[3.6]	[3.5]	[-4.3]	[-3.8]	[-2.1]	[0.5]	[-0.2]
Annual Returns			**9.5**	**9.0**	**8.1**			-0.0	0.1	0.1
			[8.8]	[4.5]	[4.8]			[-0.4]	[1.4]	[0.2]
Look Ahead	**6.0**	7.2	**8.0**	6.0	**4.9**	**-0.1**	-0.2	-0.2	0.1	**0.3**
	[12.0]	[0.9]	[4.6]	[0.5]	[4.8]	[-4.3]	[-1.3]	[-1.7]	[0.1]	[3.5]
Half Window	**5.0**	**6.4**	**7.4**	**6.7**	5.2	**-0.1**	**-0.2**	-0.1	0.1	0.2
	[7.6]	[5.2]	[3.5]	[2.6]	[0.7]	[-6.3]	[-2.7]	[-1.0]	[0.9]	[1.0]
10 Periods	**5.6**	**5.5**	**7.4**	**6.4**	1.9	**-0.1**	-0.1	-0.1	0.1	**0.4**
	[9.3]	[2.7]	[3.5]	[3.8]	[0.1]	[-3.8]	[-1.7]	[-1.0]	[1.1]	[2.1]
20 Years	**7.9**	**8.3**	**7.8**	**6.7**	5.2	**-0.2**	**-0.3**	-0.1	0.1	0.2
	[31.1]	[14.3]	[5.1]	[3.4]	[0.7]	[-5.1]	[-5.1]	[-1.9]	[0.8]	[1.0]

Table A4: Robustness of Conditional High-Frequency Approach. Data from 1927-2015 used to form non-overlapping monthly, quarterly, annual returns and annual-overlapping 3-year, 10-year returns. The table shows different specifications for calculating the conditional high-frequency approach (see Section 2 for details). *Base Case* uses daily, daily, monthly, monthly, monthly high-frequency returns with conditional means calculated using backward-looking rolling windows of 5, 10, 20, 30, 40 years for monthly, quarterly, annual, 3-year, and 10-year horizons, respectively. *No Overlap* uses non-overlapping returns (affects only 3-year and 10-year horizons which use annual overlapping horizons under the base case). *Wght: Triangular* uses as the weighting kernel for lead-lag terms a triangular (Bartlett) kernel instead of a Gaussian kernel under the base case. *Wght: Unit Closest Half* uses uniform weighting that includes only the closest half of all lead-lag terms while ignoring all others. *Cond Mean: Look Ahead* uses a forward-looking window for the conditional mean (instead of the base case's backward looking window). *Cond Mean: Half Window* uses half the window length for the conditional mean (i.e. 2.5, 5, 10, 15, 20 years, respectively). *Lower HF Frequency* uses a lower-frequency for the horizon's high-frequency returns (NA, weekly, monthly, quarterly, annually, respectively). *Den Limit* uses a different limit on the extent to which market autocorrelations can affect the beta denominator (0.5 and 0.1 limits compared to the base case of 0.3). T-statistics are shown in brackets using Newey-West '94 standard errors. Boldfaced coefficients denote a significant difference from zero at a two-sided 95% level.

(a) Size (SMB)

	Alpha					Beta				
	Mn	Qt	Yr	3Y	10Y	Mn	Qt	Yr	3Y	10Y
Base Case	**1.9**	1.0	0.6	**2.0**	**5.0**	**0.1**	**0.2**	**0.3**	0.0	**-0.3**
	[7.4]	[1.5]	[0.7]	[2.2]	[6.1]	[4.6]	[6.7]	[4.5]	[0.2]	[-2.8]
No Overlap				**1.8**	**4.7**				0.0	**-0.4**
				[2.5]	[6.1]				[0.4]	[-2.5]
Wght:	**1.8**	1.1	0.8	**2.0**	**4.9**	**0.1**	**0.2**	**0.2**	0.0	**-0.3**
Triangular	[7.7]	[1.6]	[1.0]	[2.2]	[6.3]	[4.8]	[6.1]	[4.0]	[0.2]	[-3.1]
Wght: Unit	**1.7**	0.6	0.7	**2.2**	**5.4**	**0.1**	**0.2**	**0.3**	-0.0	-0.3
Closest Half	[5.7]	[0.8]	[0.8]	[2.3]	[2.3]	[5.5]	[7.3]	[4.4]	[-0.3]	[-1.2]
Cond Mean:	**1.9**	0.7	-0.2	0.7	**2.6**	**0.1**	**0.2**	**0.3**	0.1	-0.2
Look Ahead	[5.7]	[1.1]	[-0.3]	[0.8]	[4.2]	[5.8]	[6.0]	[4.2]	[1.3]	[-1.7]
Cond Mean:	**2.0**	1.2	0.5	**1.8**	**4.9**	**0.1**	**0.2**	**0.3**	0.1	-0.1
Half Wndw	[7.0]	[1.3]	[0.5]	[2.1]	[4.5]	[4.8]	[6.6]	[5.2]	[1.5]	[-0.5]
Lower HF		1.1	0.6	**1.9**	**5.1**		**0.2**	**0.3**	0.0	**-0.3**
Frequency		[1.6]	[0.7]	[2.2]	[5.3]		[5.7]	[4.5]	[0.4]	[-2.3]
Den Limit:	**1.9**	1.1	0.6	**1.9**	**4.8**	**0.1**	**0.2**	**0.2**	0.0	**-0.3**
0.5	[8.5]	[1.7]	[0.7]	[2.1]	[4.6]	[4.9]	[6.4]	[4.7]	[0.3]	[-2.3]
Den Limit:	**1.9**	0.9	0.6	**1.9**	**5.2**	**0.1**	**0.2**	**0.3**	0.0	**-0.4**
0.1	[6.9]	[1.3]	[0.7]	[2.3]	[6.2]	[4.5]	[6.5]	[4.4]	[0.3]	[-3.1]

(b) Value (HML)

	Alpha					Beta				
	Mn	Qt	Yr	3Y	10Y	Mn	Qt	Yr	3Y	10Y
Base Case	**4.7**	**4.3**	**3.6**	**4.4**	**5.2**	-0.0	0.0	0.1	-0.0	-0.1
	[19.3]	[10.0]	[4.3]	[7.9]	[6.9]	[-0.5]	[0.2]	[0.6]	[-0.1]	[-0.4]
No Overlap				**3.4**	**4.6**				0.1	-0.0
				[3.0]	[11.4]				[0.7]	[-0.0]
Wght:	**4.7**	**4.3**	**3.7**	**4.4**	**5.3**	-0.0	0.0	0.0	-0.0	-0.1
Triangular	[18.4]	[9.8]	[4.8]	[8.9]	[7.1]	[-0.4]	[0.2]	[0.5]	[-0.2]	[-0.6]
Wght: Unit	**4.7**	**4.1**	**3.4**	**4.4**	**6.0**	-0.0	0.0	0.1	-0.0	-0.1
Closest Half	[17.6]	[8.8]	[3.5]	[7.3]	[2.3]	[-0.6]	[0.9]	[0.7]	[-0.1]	[-0.4]
Cond Mean:	**3.1**	**2.8**	**2.8**	3.0	**3.6**	-0.0	0.0	0.0	-0.0	-0.0
Look Ahead	[8.5]	[4.3]	[3.7]	[1.9]	[4.1]	[-0.2]	[0.3]	[0.3]	[-0.1]	[-0.1]
Cond Mean:	**4.6**	**4.4**	**3.8**	**3.9**	**5.1**	-0.0	0.0	0.0	0.0	-0.1
Half Wndw	[16.3]	[6.3]	[4.0]	[6.6]	[6.7]	[-0.5]	[0.2]	[0.3]	[0.1]	[-0.8]
Lower HF		**4.3**	**3.6**	**4.4**	**5.4**		0.0	0.1	-0.0	-0.1
Frequency		[10.4]	[4.3]	[7.7]	[6.6]		[0.1]	[0.6]	[-0.0]	[-0.5]
Den Limit:	**4.6**	**4.2**	**3.6**	**4.4**	**5.2**	-0.0	0.0	0.0	0.0	-0.1
0.5	[19.8]	[9.8]	[5.0]	[8.8]	[6.7]	[-0.5]	[0.2]	[0.5]	[0.0]	[-0.4]
Den Limit:	**4.7**	**4.2**	**3.5**	**4.3**	**5.0**	-0.0	0.0	0.1	-0.0	-0.0
0.1	[18.1]	[10.2]	[3.9]	[7.0]	[8.3]	[-0.3]	[0.5]	[0.8]	[-0.1]	[-0.2]

(c) Momentum (UMD)

	Alpha					Beta				
	Mn	Qt	Yr	3Y	10Y	Mn	Qt	Yr	3Y	10Y
Base Case	**6.0**	**7.0**	**7.2**	**7.4**	4.5	**-0.0**	-0.1	-0.1	-0.0	0.2
	[12.1]	[7.1]	[6.3]	[6.0]	[1.3]	[-2.1]	[-1.0]	[-0.7]	[-0.6]	[1.1]
No Overlap				**6.7**	**6.2**				0.0	0.1
				[5.5]	[3.3]				[0.5]	[0.7]
Wght:	**5.9**	**6.8**	**7.3**	**7.3**	4.5	-0.0	-0.0	-0.1	-0.0	0.2
Triangular	[12.1]	[6.9]	[6.4]	[5.9]	[1.2]	[-1.5]	[-0.6]	[-0.9]	[-0.6]	[1.1]
Wght: Unit	**6.2**	**7.3**	**7.4**	**7.7**	2.9	**-0.1**	-0.1	-0.1	-0.0	0.4
Closest Half	[11.8]	[7.4]	[6.2]	[6.4]	[0.7]	[-2.7]	[-1.7]	[-0.9]	[-0.2]	[1.3]
Cond Mean:	**7.8**	**7.2**	**7.5**	**6.4**	**5.4**	-0.0	-0.1	-0.1	0.0	0.2
Look Ahead	[20.7]	[10.0]	[8.9]	[8.3]	[3.2]	[-2.0]	[-0.9]	[-0.8]	[0.1]	[1.7]
Cond Mean:	**4.6**	**6.1**	**6.7**	**7.5**	5.3	**-0.1**	-0.1	-0.1	-0.0	0.2
Half Wndw	[8.3]	[5.1]	[5.0]	[4.5]	[1.5]	[-2.5]	[-1.0]	[-1.2]	[-0.5]	[1.9]
Lower HF		**7.0**	**7.2**	**7.4**	4.3		-0.1	-0.1	-0.1	0.2
Frequency		[7.4]	[6.3]	[6.1]	[1.2]		[-1.3]	[-0.7]	[-0.8]	[1.2]
Den Limit:	**6.0**	**7.0**	**7.4**	**7.4**	4.3	-0.0	-0.1	-0.1	-0.0	0.3
0.5	[13.0]	[7.4]	[6.5]	[6.1]	[1.0]	[-1.9]	[-1.1]	[-0.8]	[-0.7]	[1.2]
Den Limit:	**6.0**	**6.9**	**7.0**	**7.5**	4.8	**-0.0**	-0.0	-0.0	-0.1	0.2
0.1	[11.8]	[6.8]	[5.8]	[5.6]	[1.4]	[-2.1]	[-0.7]	[-0.4]	[-0.7]	[0.8]

References

Allaire, JJ, Joe Cheng, Yihui Xie, Jonathan McPherson, Winston Chang, Jeff Allen, Hadley Wickham, Aron Atkins, and Rob Hyndman. 2016. *Rmarkdown: Dynamic Documents for R*. https://CRAN.R-project.org/package=rmarkdown.

Ameriks, John, and Stephen P Zeldes. 2004. "How Do Household Portfolio Shares Vary with Age." *Working Paper*.

Andersen, Torben, and Tim Bollerslev. 1998. "Answering the Skeptics: Yes, Standard Volatility Models Do Provide Accurate Forecasts." *International Economic Review*. JSTOR, 885–905.

Andersen, Torben, Tim Bollerslev, and Francis Diebold. 2010. "Parametric and Nonparametric Measurements of Volatility." In *Handbook of Financial Econometrics*, edited by Yacine Ait-Sahalia and Lars Peter Hansen, 67–137. Elsevier.

Andersen, Torben, Tim Bollerslev, Francis Diebold, and Jin Wu. 2006. "Realized Beta: Persistence and Predictability." *Advances in Econometrics* 20 (2). JAI Press: 1–39.

Andries, Marianne, Thomas M Eisenbach, Martin C Schmalz, and Yichuan Wang. 2015. "The Term Structure of the Price of Variance Risk." *FRB of New York Working Paper No. FEDNSR736*.

Andries, Marianne, Thomas Eisenbach, and Martin Schmalz. 2015. "Asset Pricing with Horizon-Dependent Risk Aversion." *FRB of New York Staff Report*, no. 703.

Ang, Andrew, and Dennis Kristensen. 2012. "Testing Conditional Factor Models." *Journal of Financial Economics* 106 (1). Elsevier: 132–56.

Ang, Andrew, and Jun Liu. 2004. "How to Discount Cashflows with Time-Varying Expected Returns." *The Journal of Finance* 59 (6). Wiley Online Library: 2745–83.

Ang, Andrew, and Maxim Ulrich. 2012. "Nominal Bonds, Real Bonds, and Equity." Netspar Discussion Paper.

Ang, Andrew, Geert Bekaert, and Min Wei. 2008. "The Term Structure of Real Rates and Expected Inflation." *The Journal of Finance* 63 (2). Wiley Online Library: 797–849.

Asness, Clifford, Andrea Frazzini, Lasse Pedersen, Ronen Israel, and Tobias Moskowitz. 2015. "Size Matters, If You Control Your Junk." *Working Paper*.

Asparouhova, Elena, Hendrik Bessembinder, and Ivalina Kalcheva. 2013. "Noisy Prices and Inference Regarding Returns." *The Journal of Finance* 68 (2). Wiley Online Library: 665–714.

Bache, Stefan Milton, and Hadley Wickham. 2014. *Magrittr: A Forward-Pipe Operator for R*. https://CRAN.R-project.org/package=magrittr.

Bandi, Federico M, and Benoit Perron. 2008. "Long-Run Risk-Return Trade-Offs." *Journal of Econometrics* 143 (2). Elsevier: 349–74.

Bandi, Federico M, Renè Garcia, Abraham Lioui, and Benoît Perron. 2010. "A Long Horizon Perspective on the Cross Section of Expected Returns." *Working Paper*.

Barndorff-Nielsen, Ole, and Neil Shephard. 2004. "Econometric Analysis of Realized Covariation: High Frequency Based Covariance, Regression, and Correlation in Financial Economics."

Econometrica 72 (3). Wiley Online Library: 885–925.

Barroso, Pedro, and Pedro Santa-Clara. 2015. "Momentum Has Its Moments." *Journal of Financial Economics* 116 (1). Elsevier: 111–20.

Belo, Frederico, Pierre Collin-Dufresne, and Robert Goldstein. 2015. "Dividend Dynamics and the Term Structure of Dividend Strips." *The Journal of Finance* 70 (3). Wiley Online Library: 1115–60.

Boguth, Oliver, Murray Carlson, Adlai Fisher, and Mikhail Simutin. 2016. "Horizon Effects in Average Returns: The Role of Slow Information Diffusion." *Review of Financial Studies*. Soc Financial Studies, hhw024.

Boons, Martijn, and Andrea Tamoni. 2015. "Horizon-Specific Macroeconomic Risks and the Cross-Section of Expected Returns." *Available at SSRN 2516251*.

Brennan, Michael J, and Yuzhao Zhang. 2013. "Capital Asset Pricing with a Stochastic Horizon." *Available at SSRN 1752654*.

Bryzgalova, Svetlana, and Christian Julliard. 2015. "The Consumption Risk of Bonds and Stocks."

Campbell, John, and Luis Viceira. 2002. *Strategic Asset Allocation: Portfolio Choice for Long-Term Investors*. Oxford University Press.

———. 2005. "The Term Structure of the Risk-Return Trade-Off." *Financial Analysts Journal* 61 (1). CFA Institute: 34–44.

Cella, Cristina, Andrew Ellul, and Mariassunta Giannetti. 2013. "Investors' Horizons and the Amplification of Market Shocks." *Review of Financial Studies*. Soc Financial Studies, hht023.

Chaudhuri, Shomesh E, and Andrew W Lo. 2015. "Spectral Analysis of Stock-Return Volatility, Correlation, and Beta." In *Signal Processing and Signal Processing Education Workshop (Sp/Spe), 2015 Ieee*, 232–36. IEEE.

Chausse, Pierre. 2010. "Computing Generalized Method of Moments and Generalized Empirical Likelihood with R." *Journal of Statistical Software* 34 (11): 1–35. http://www.jstatsoft.org/v34/i11/.

Chordia, Tarun, and Bhaskaran Swaminathan. 2000. "Trading Volume and Cross-Autocorrelations in Stock Returns." *The Journal of Finance* 55 (2). Wiley Online Library: 913–35.

Chordia, Tarun, Asani Sarkar, and Avanidhar Subrahmanyam. 2011. "Liquidity Dynamics and Cross-Autocorrelations." *Journal of Financial and Quantitative Analysis* 46 (03). Cambridge Univ Press: 709–36.

Cohen, Kalman, Gabriel Hawawini, Steven Maier, Robert Schwartz, and David Whitcomb. 1983. "Friction in the Trading Process and the Estimation of Systematic Risk." *Journal of Financial Economics* 12 (2). Elsevier: 263–78.

Cohen, Lauren, and Andrea Frazzini. 2008. "Economic Links and Predictable Returns." *The Journal of Finance* 63 (4). Wiley Online Library: 1977–2011.

Cohen, Lauren, and Dong Lou. 2012. "Complicated Firms." *Journal of Financial Economics* 104 (2). Elsevier: 383–400.

Cohen, Randolph, Christopher Polk, and Tuomo Vuolteenaho. 2003. "The Value Spread." *The Journal of Finance* 58 (2). Wiley Online Library: 609–42.

———. 2009. "The Price Is (Almost) Right." *The Journal of Finance* 64 (6). Wiley Online Library:

2739–82.

Colacito, Riccardo, and Robert Engle. 2010. "Term Structure of Risk, the Role of Known and Unknown Risks and Non-Stationary Distributions."

Conrad, Jennifer, and Gautam Kaul. 1989. "Mean Reversion in Short-Horizon Expected Returns." *Review of Financial Studies* 2 (2). Soc Financial Studies: 225–40.

Croce, Mariano M, Martin Lettau, and Sydney C Ludvigson. 2014. "Investor Information, Long-Run Risk, and the Term Structure of Equity." *Review of Financial Studies*. Soc Financial Studies, hhu084.

Daniel, Kent, and Tobias Moskowitz. 2015. "Momentum Crashes." *Working Paper*.

Dew-Becker, Ian, and Stefano Giglio. 2016. "Asset Pricing in the Frequency Domain: Theory and Empirics." *Review of Financial Studies*. Soc Financial Studies, hhw027.

Dimson, Elroy. 1979. "Risk Measurement When Shares Are Subject to Infrequent Trading." *Journal of Financial Economics* 7 (2). Elsevier: 197–226.

Diris, Bart, Franz Palm, and Peter Schotman. 2014. "Long-Term Strategic Asset Allocation: An Out-of-Sample Evaluation." *Management Science* 61 (9). INFORMS: 2185–2202.

Dowle, M, A Srinivasan, T Short, and S Lianoglou. 2015. *Data.table: Extension of Data.frame*. https://CRAN.R-project.org/package=data.table.

Fama, Eugene, and Kenneth French. 1988. "Permanent and Temporary Components of Stock Prices." *The Journal of Political Economy*. JSTOR, 246–73.

———. 1989. "Business Conditions and Expected Returns on Stocks and Bonds." *Journal of Financial Economics* 25 (1). Elsevier: 23–49.

———. 1993. "Common Risk Factors in the Returns on Stocks and Bonds." *Journal of Financial Economics* 33 (1). Elsevier: 3–56.

———. 2012. "Size, Value, and Momentum in International Stock Returns." *Journal of Financial Economics* 105 (3). Elsevier: 457–72.

Fama, Eugene, and James MacBeth. 1973. "Risk, Return, and Equilibrium: Empirical Tests." *The Journal of Political Economy*. JSTOR, 607–36.

Ferson, Wayne, and Campbell Harvey. 1999. "Conditioning Variables and the Cross Section of Stock Returns." *The Journal of Finance* 54 (4). Wiley Online Library: 1325–60.

Ferson, Wayne, and Rudi Schadt. 1996. "Measuring Fund Strategy and Performance in Changing Economic Conditions." *The Journal of Finance* 51 (2). Wiley Online Library: 425–61.

Fleming, Jeff, Chris Kirby, and Barbara Ostdiek. 2003. "The Economic Value of Volatility Timing Using Realized Volatility." *Journal of Financial Economics* 67 (3). Elsevier: 473–509.

Fowler, David, and Harvey Rorke. 1983. "Risk Measurement When Shares Are Subject to Infrequent Trading: Comment." *Journal of Financial Economics* 12 (2). Elsevier: 279–83.

Gencay, Ramazan, Faruk Selcuk, and Brandon Whitcher. 2005. "Multiscale Systematic Risk." *Journal of International Money and Finance* 24 (1). Elsevier: 55–70.

Gilbert, Thomas, Christopher Hrdlicka, Jonathan Kalodimos, and Stephan Siegel. 2014. "Daily

Data Is Bad for Beta: Opacity and Frequency-Dependent Betas.” *Review of Asset Pricing Studies* 4 (1). Oxford University Press: 78–117.

Goyal, Amit, and Ivo Welch. 2008. “A Comprehensive Look at the Empirical Performance of Equity Premium Prediction.” *Review of Financial Studies* 21 (4). Soc Financial Studies: 1455–1508.

Guidolin, Massimo, and Allan Timmermann. 2006. “Term Structure of Risk Under Alternative Econometric Specifications.” *Journal of Econometrics* 131 (1). Elsevier: 285–308.

Handa, Puneet, Sagar Kothari, and Charles Wasley. 1989. “The Relation Between the Return Interval and Betas: Implications for the Size Effect.” *Journal of Financial Economics* 23 (1). Elsevier: 79–100.

Handa, Puneet, SP Kothari, and Charles Wasley. 1993. “Sensitivity of Multivariate Tests of the Capital Asset-Pricing Model to the Return Measurement Interval.” *The Journal of Finance* 48 (4). Wiley Online Library: 1543–51.

Harvey, Campbell, Yan Liu, and Heqing Zhu. 2015.“. And the Cross-Section of Expected Returns.” *Review of Financial Studies*. Soc Financial Studies, hhv059.

Hawawini, Gabriel. 1983. “Why Beta Shifts as the Return Interval Changes.” *Financial Analysts Journal* 39 (3). CFA Institute: 73–77.

Hong, Harrison, Walter Torous, and Rossen Valkanov. 2007. “Do Industries Lead Stock Markets?” *Journal of Financial Economics* 83 (2). Elsevier: 367–96.

Hou, Kewei. 2007. “Industry Information Diffusion and the Lead-Lag Effect in Stock Returns.” *Review of Financial Studies* 20 (4). Soc Financial Studies: 1113–38.

Hou, Kewei, and Mathijs Van Dijk. 2014. “Resurrecting the Size Effect: Firm Size, Profitability Shocks, and Expected Stock Returns.” *Charles A. Dice Center Working Paper*, no. 2010-1.

In, Francis, Sangbae Kim, and Robert Faff. 2010. “Explaining Mispricing with Fama–French Factors: New Evidence from the Multiscaling Approach.” *Applied Financial Economics* 20 (4). Taylor & Francis: 323–30.

In, Francis, Sangbae Kim, and Ramazan Gencay. 2011. “Investment Horizon Effect on Asset Allocation Between Value and Growth Strategies.” *Economic Modelling* 28 (4). Elsevier: 1489–97.

Jagannathan, Ravi, and Zhenyu Wang. 1996. “The Conditional Capm and the Cross-Section of Expected Returns.” *The Journal of Finance* 51 (1). Wiley Online Library: 3–53.

Jurek, Jakub, and Luis Viceira. 2011. “Optimal Value and Growth Tilts in Long-Horizon Portfolios.” *Review of Finance* 15 (1). Oxford University Press: 29–74.

Kamara, Avraham, Robert Korajczyk, Xiaoxia Lou, and Ronnie Sadka. 2015. “Horizon Pricing.” *Working Paper*.

Kendall, Maurice. 1954. “Note on Bias in the Estimation of Autocorrelation.” *Biometrika* 41 (3-4). Biometrika Trust: 403–4.

Levhari, David, and Haim Levy. 1977. “The Capital Asset Pricing Model and the Investment Horizon.” *The Review of Economics and Statistics*. JSTOR, 92–104.

Levi, Yaron, and Ivo Welch. 2016. “Best Practice for Cost-of-Capital Estimates.” *Available at*

SSRN 2667735.

Lewellen, Jonathan, and Stefan Nagel. 2006. "The Conditional Capm Does Not Explain Asset-Pricing Anomalies." *Journal of Financial Economics* 82 (2). Elsevier: 289–314.

Liu, Weimin, and Norman Strong. 2008. "Biases in Decomposing Holding-Period Portfolio Returns." *Review of Financial Studies* 21 (5). Soc Financial Studies: 2243–74.

Lo, Andrew, and Craig MacKinlay. 1990. "When Are Contrarian Profits Due to Stock Market Overreaction?" *Review of Financial Studies* 3 (2). Soc Financial Studies: 175–205.

Menzly, Lior, and Oguzhan Ozbas. 2010. "Market Segmentation and Cross-Predictability of Returns." *The Journal of Finance* 65 (4). Wiley Online Library: 1555–80.

Merton, Robert. 1969. "Lifetime Portfolio Selection Under Uncertainty: The Continuous-Time Case." *The Review of Economics and Statistics*. JSTOR, 247–57.

———. 1980. "On Estimating the Expected Return on the Market: An Exploratory Investigation." *Journal of Financial Economics* 8 (4). Elsevier: 323–61.

Novy-Marx, Robert, and Mihail Velikov. 2016. "A Taxonomy of Anomalies and Their Trading Costs." *Review of Financial Studies* 29 (1). Soc Financial Studies: 104–47.

Poterba, James, and Lawrence Summers. 1988. "Mean Reversion in Stock Prices: Evidence and Implications." *Journal of Financial Economics* 22 (1). Elsevier: 27–59.

Rua, Antonio, and Luis C Nunes. 2012. "A Wavelet-Based Assessment of Market Risk: The Emerging Markets Case." *The Quarterly Review of Economics and Finance* 52 (1). Elsevier: 84–92.

Samuelson, Paul. 1969. "Lifetime Portfolio Selection by Dynamic Stochastic Programming." *The Review of Economics and Statistics*. JSTOR, 239–46.

Scholes, Myron, and Joseph Williams. 1977. "Estimating Betas from Nonsynchronous Data." *Journal of Financial Economics* 5 (3). Elsevier: 309–27.

Van Binsbergen, Jules, Michael Brandt, and Ralph Koijen. 2012. "On the Timing and Pricing of Dividends." *American Economic Review* 102 (4). American Economic Association: 1596–1618.

Wickham, Hadley. 2016. *Tidyr: Easily Tidy Data with 'Spread()' and 'Gather()' Functions*. https://CRAN.R-project.org/package=tidyr.

Wickham, Hadley, and Romain Francois. 2016. *Dplyr: A Grammar of Data Manipulation*. https://CRAN.R-project.org/package=dplyr.

Zeileis, Achim. 2004. "Econometric Computing with Hc and Hac Covariance Matrix Estimators." *Journal of Statistical Software*. http://www.jstatsoft.org/v34/i11/.

Zeileis, Achim, and Gabor Grothendieck. 2005. "Zoo: S3 Infrastructure for Regular and Irregular Time Series." *Journal of Statistical Software* 14 (6): 1–27. doi:10.18637/jss.v014.i06.