Forecasting Realized Betas Using Predictors Indicating Structural Breaks and Asymmetric Risk Effects

Jiawen Luo[a, *1], Zhenbiao Chen[a], Mingmian Cheng[b]

[a] *School of Business Administration, South China University of Technology, Guangzhou 510640, China*

[b] *Lingnan College, Sun Yat-sen University, Guangzhou 510275, China*

Abstract: This paper studies the importance of structural breaks and asymmetric risk effects for accurate forecasts of the realized beta. Specifically, structural breaks in the realized beta are detected by Iterated Cumulative Sum of Square (ICSS) algorithm and asymmetric risk effects are captured by decomposing the realized beta further into various components following Ang et al. (2006) and Bollerslev et al. (2021). We propose a set of Heterogeneous Autoregressive (HAR) model variants by incorporating these new predictors. To achieve model parsimony and to keep only the predictors with significant power, we employ Least Absolute Shrinkage and Selection Operator (LASSO) method for variable selection. Our proposed LASSO-HAR model with estimators of structural breaks and asymmetric risk effects is found to yield more accurate out-of-sample beta forecasts than a variety of alternative models in terms of both statistical and economic criteria. In particular, our model successfully achieves the long-memory feature of realized betas in a tractable and parsimonious way. These empirical findings are robust across different data sampling frequencies, different estimation windows, different sub-samples, different quantiles of the beta distribution and different industrial sectors.

Keywords: Realized beta; structural break; asymmetric risk; HAR model; LASSO

JEL Classification: C22, C52, C53, G17

---

[1] * Corresponding author. School of Business Administration, South China University of Technology, Guangzhou, China. Email: 2jialuo@163.com.

1. Introduction

Capital Asset Pricing Model (CAPM) (Sharpe, 1964; Linter, 1965) is one of the most important theoretical models in modern finance, which has a wide range of applications in asset pricing, portfolio allocation and risk management. CAPM betas (or factor loadings) are interpretable as exposure to systematic risk factors that drive expected returns. However, one of the main obstacles to the empirical analysis of factor models is that betas are not directly observable. As a result, it is of great theoretical and practical significance to accurately estimating and forecasting betas (Tofallis, 2008; Levi and Welch, 2017; Daniel et al., 2020).

In earlier research, monthly beta estimators constructed using daily asset price data were widely used. Although these estimators are consistent under weak regulatory conditions, they fail to fully explain the significant long-memory feature that is universally observed in the time series of betas (Baillie et al., 1996; Bollerslev and Mikkelsen, 1996; Ding and Granger, 1996). Andersen et al. (2006) firstly proposed the realized beta estimator that is allowed to change over time using intra-daily asset prices. With the ever-increasing availability of high-frequency financial data, the realized beta provides very accurate estimates and, not surprisingly, has become one of the most popular measures for betas in both academia and industry. Wang (2003), Ghysels and Jacquier (2006), and many others emphasized the importance of accurate beta forecasts for the successful construction of market-neutral portfolios. In this regard, numerous researchers investigated various time-series models for forecasting realized betas. For instance, Hooper et al. (2008) forecast quarterly realized betas and found that the autoregressive model with two lags far surpassed the constant beta model; Reeves and Wu (2013) evaluated the forecasting performance of the constant beta model against autoregressive models for time-varying betas; Cenesizoglu et al. (2017) forecast realized betas with a variety of autoregressive models over long forecast horizons and evaluated the statistical accuracy and economic significance of different forecasting approaches; More recently, Drobetz et al. (2021) conducted a comprehensive study of beta forecasting across a large number of developed and emerging markets. They found that the forecasting approaches with estimators constructed using daily data outperformed those using monthly or quarterly data. Most of the forecasting models considered in the aforementioned papers and the references cited therein are short-memory models. However, Becker et al. (2021) provided the empirical evidence suggesting that monthly realized betas exhibit consistent characteristics of long-run dependence.

Therefore, they investigated the true long-memory models and documented superior forecasting performance, when compared to short-memory and difference-stationary forecasting models.

Besides the long-memory feature, the asymmetric risk effects of realized betas have also been widely discussed in the extant literature. Ang et al. (2006) indicated that realized betas can be divided into upside and downside components. They also showed that the CAPM model with downside betas can better account for the cross-sectional variation in U.S. equity returns. Post and van Vliet (2004), Lettau et al., (2014) and Zabarankin et al. (2014) reached the same conclusion for other asset classes. By contrast, a recent work by Atilgan et al. (2018) called into question the ability of downside betas to successfully explain the cross-sectional variation in U.S. and international equity returns. In addition, the predictability of different beta components has also been studied. For instance, Levi and Welch (2020) suggested that betas contain more relevant information for forecasting downside betas because downside betas were more prone to measurement errors. Recently, Bollerslev et al. (2021) proposed to decompose realized betas into four semi-beta components and explored the explanatory power of these realized semi-betas on asset pricing. These estimators, as well as upside and downside betas, that proxy for various asymmetric risk effects may provide additional information for forecasting realized betas.

To account for the empirical features of realized betas (e.g., long memory, asymmetric risk effects, and structural breaks), we propose a set of Heterogeneous Autoregressive (HAR) model variants. Specifically, the long-memory properties are achieved by the additive beta cascade of HAR model (Corsi, 2009); The structural breaks are detected by Iterated Cumulative Sum of Square (ICSS) algorithm and the structural break dummies are subsequently added as predictors; The asymmetric risk effects are quantified by different beta components, including the four types of semi-betas and the upside and downside betas. However, incorporating all these predictors into the HAR model may potentially lead to overfitting issues. To keep model parsimony and to identify the predictors with significant power, we employ the Least Absolute Shrinkage and Selection Operator (LASSO) method for variable selection. Much of the extant literature has showed that the LASSO method can lead to lower risk of misidentification for forecasting models and thus better out-of-sample forecasting accuracy (Siliverstovs, 2015; Ziel and Liu, 2016; Sagaert et al., 2018; Liang et al., 2022).

We contribute to the literature in at least three aspects. First, we focus on the forecasts of daily realized betas estimated by using five-minute high-frequency data. Papageorgiou et al. (2016) showed that it was hard to construct market-neutral portfolios with traditional beta forecasts and suggested the need of high-frequency data. Second, we propose a set of new models accounting for long memory, structural breaks and asymmetric risk effects. Compared with the true long-memory models employed in Becker et al. (2021), our proposed models are much more parsimonious and easier to implement. The empirical results suggest that the forecasting performance is substantially improved in term of both statistical and economic criteria. We also find that both upside and downside betas are useful for risk management and hedging purposes. Third, we further introduce the LASSO method for variable selection and capture the dynamic changes of the financial market. The use of LASSO method not only selects the powerful predictors but also improves the predictive accuracy and economic value of the HAR-type models.

The remainder of this paper is organized as follows. Section 2 introduces the high-frequency data sets and statistical methods for the construction of various estimators. Section 3 describes the forecasting models. Section 4 documents the out-of-sample forecasting results and provides the economic value analysis. Section 5 presents further analysis results. Section 6 demonstrates the robustness check results. Section 7 concludes.

2. Data and Statistical Methodology

2.1. Data

We collect intra-daily price data (Hollstein et al., 2020; Becker et al., 2021) sampled at a five-minute frequency on 327 constituents of S&P 500 index from the Trade and Quote (TAQ) database,[2] for the full sample period from January 2007 to December 2019, covering 3066 trading days in total. The stocks that are collected in our data set account for 63.21% of the entire market capitalization of S&P 500 index in December 2019. Overnight price data are excluded from our data set, i.e., we use only data with time stamps between 9:30AM and 4:00PM Eastern Standard Time.

[2] As the constituents of S&P 500 index have been changing over time, only stocks that are present on all trading days from 2007 to 2019 are kept. In addition, stocks with missing data are dropped. As a result, there are in total 327 stocks in our data set.

2.2. Realized beta

A consistent estimator of beta is required since it is unobservable. Historically, betas were estimated by using low-frequency data. However, Andersen et al. (2005) indicated that such an approach can yield noisy and thus inaccurate estimates. Taking advantage of high-frequency data, Andersen et al. (2006) firstly proposed the realized beta that has been attracting ever-increasing attention. On the one hand, the realized beta is theoretically proven to be a consistent estimator of the integrated beta. On the other hand, much empirical support has been found for the use of this realized measure in the extant literature. For instance, Patton and Verardo (2012) argued that detecting variations in betas measured at higher frequencies is crucial to understanding the effect of information flows on the covariance structure of stock returns. As a result, the variable of primary interest in this paper is chosen as the realized beta that serves as a proxy for the true unobservable beta. When constructing realized betas from high-frequency data, the impact of microstructure noise may severely deteriorate the accuracy of beta estimates (Bhattacharyya et al., 2009). Due to such concerns, we adopt the five-minute sampling scheme for the estimation of daily realized betas.[3] Let $r_{i,\tau}$ and $r_{M,\tau}$ denote, respectively, the intra-daily high-frequency returns on asset *i* and on the aggregate market over the τ^{th} time interval, the daily realized beta of asset *i* on day *t* can be constructed as follows,

$$\beta_{i,t} = \frac{\sum_{\tau=1}^{O} r_{i,\tau} r_{M,\tau}}{\sum_{\tau=1}^{O} r_{M,\tau}^2}, (i = 1,2,\ldots,N;\ t = 1,2,\ldots,T), \quad (1)$$

where *O* is the total number of high-frequency return observations within day *t*, *N* is the number of individual stocks, and *T* refers to the full-time span of our data set.

To validate the choice of sampling frequency, we also rely on the signature plot method of Andersen et al. (2000), as illustrated in Figure 1. It can be observed that the average realized betas increase steadily from $k = 50$ (corresponding to a 50-minute return sampling interval) with a value of 0.835. When using five-minute returns to construct daily realized betas, the average realized beta is 0.832, which is approximately equal to the value at $k = 50$. Therefore, selecting a five-minute interval for constructing realized beta estimators strikes a reasonable balance between minimizing microstructure noise and reducing estimation bias.

[3] As a robustness check, we also constructed daily and monthly realized betas using 30-minute high-frequency data. The empirical findings are qualitatively similar to those presented in our paper.

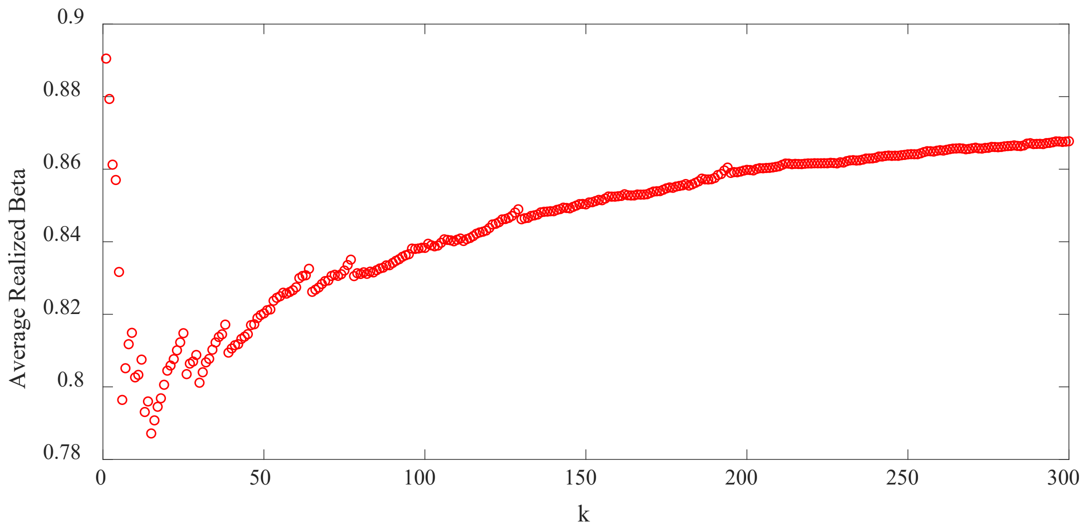

Figure 1. Realized betas signature plot. This figure shows the average realized betas for daily realized betas across all 327 individual stocks, using various *k*-minute return intervals.

Much empirical evidence has shown that betas consistently exhibit significant long-memory properties. To examine this main empirical feature of realized betas, we employ the Geweke and Porter-Hudak (GPH, 1983) method to estimate the fractional differencing parameter d for the time series of daily realized betas. The GPH estimator is unbiased for the memory parameter d, which can be applied to distinguish short-memory series ($-0.5< d < 0$), stationary long-memory series ($0 < d < 0.5$), nonstationary long-memory series ($0.5 < d <1$), and difference-stationary series ($d = 1$). We estimate the parameter d for each of the 327 individual stocks. The sample average and standard deviation of d values are reported in Table 1. On average, the daily realized beta series shows a memory parameter $d = 0.451$. Figure 2 displays the empirical distribution of all memory parameter estimates. It is observed that almost all of the d estimates are between 0 and 1. We also conduct a block bootstrap test for the null hypotheses of $d = 0$ and $d = 1$ at the 10% significance level. For approximately 90.3% and 99.6% of the individual stocks, we reject the null hypotheses of $d = 0$ and $d = 1$, respectively. In a nutshell, the time series of realized betas should be characterized as a long-memory process. Furthermore, a majority of individual stocks have stationary long-memory betas with $0 < d < 0.5$.

Table 1

Long-memory feature of realized betas. This table reports the average $\bar{\hat{d}}_i$ and the standard deviation $sd(\hat{d}_i)$ for daily realized betas across all 327 individual stocks. The memory parameter d is estimated using the GPH method. The frequencies at which we reject the null hypotheses of d_i= 0 and d_i = 1 at the 10% level, based on the block bootstrap method, are also reported.

	$\bar{\hat{d}}_i$	$sd(\hat{d}_i)$	vs. d_i= 0	vs. d_i= 1
β_i	0.451	0.105	0.903	0.996

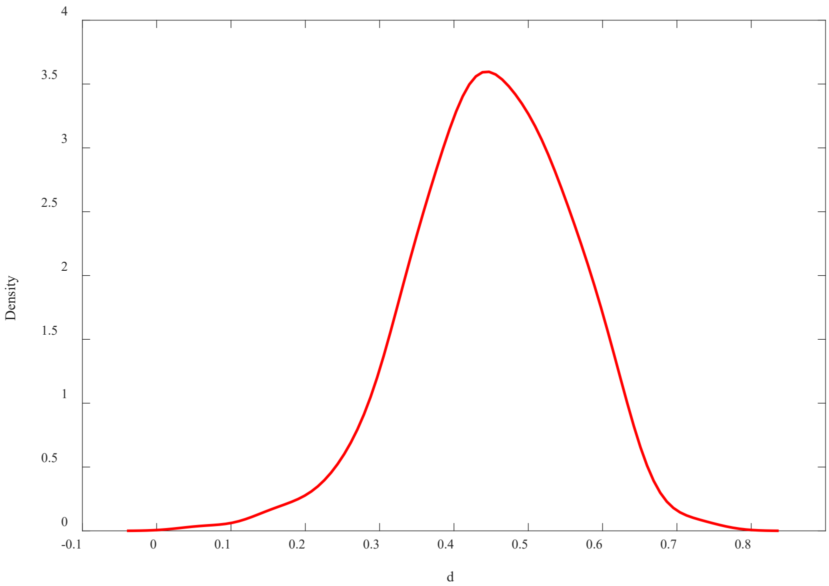

Figure 2. Empirical distribution of *d* estimates. This figure shows the empirical distribution of the estimated memory parameters for daily realized betas across all 327 individual stocks.

2.3. Realized semi-betas, realized upside and downside betas

Investor's aversion to risk leads to the creation of "betting against beta" investment strategy (Frazzini and Pedersen, 2014), particularly averse to downside risk. To exploit the asymmetric dependencies between individual stocks and the aggregate market, Bollerslev et al. (2021) decomposed the realized beta into four realized semi-beta components as follow,

$$\beta_{i,t}^{P} = \frac{\sum_{\tau=1}^{O} r_{i,\tau}^{+} r_{M,\tau}^{+}}{\sum_{\tau=1}^{O} r_{M,\tau}^{2}}, \beta_{i,t}^{N} = \frac{\sum_{\tau=1}^{O} r_{i,\tau}^{-} r_{M,\tau}^{-}}{\sum_{\tau=1}^{O} r_{M,\tau}^{2}}, \beta_{i,t}^{M+} = \frac{-\sum_{\tau=1}^{O} r_{i,\tau}^{-} r_{M,\tau}^{+}}{\sum_{\tau=1}^{O} r_{M,\tau}^{2}}, \beta_{i,t}^{M-} = \frac{-\sum_{\tau=1}^{O} r_{i,\tau}^{+} r_{M,\tau}^{-}}{\sum_{\tau=1}^{O} r_{M,\tau}^{2}}, \tag{2}$$

This decomposition relies on the new semi-covariance concept for decomposing the systematic market risk. As a result, realized semi-betas contain fundamentally different information from those defined based on asset-specific "good" and "bad" volatility measures. The realized semi-beta estimators provide an exact four-way decomposition of the classic realized beta, which is shown in (3),

$$\beta_{i,t} = \frac{\sum_{\tau=1}^{O} r_{i,\tau} r_{M,\tau}}{\sum_{\tau=1}^{O} r_{M,\tau}^{2}} = \beta_{i,t}^{P} + \beta_{i,t}^{N} - \beta_{i,t}^{M+} - \beta_{i,t}^{M-}, \tag{3}$$

If the market and individual asset returns were jointly normally distributed, it can be shown that $\beta_i^P = \beta_i^N \neq 1/2\pi(\sqrt{\sigma_i^2/\sigma_M^2 - {\beta_i}^2} + \beta_i \arccos(-\sigma_i\beta_i/\sigma_M))$ and $\beta_i^{M+} = \beta_i^{M-} \neq 1/2\pi(\sqrt{\sigma_i^2/\sigma_M^2 - {\beta_i}^2} + \beta_i \arccos(\sigma_i\beta_i/\sigma_M))$ as $O \to \infty$ (see Bollerslev et al., 2020). Otherwise, the four realized semi-betas would generally differ, which convey additional useful information to that of the standard market beta.

Apart from the four realized semi-beta measures, Ang et al. (2006) introduced upside and downside betas which were found to be useful for improving upon the traditional CAPM. The realized version of upside and downside betas can thus be defined, respectively, as follow,

$$\beta_{i,t}^{+} = \frac{\sum_{\tau=1}^{O} r_{i,\tau} r_{M,\tau}^{+}}{\sum_{\tau=1}^{O} (r_{M,\tau}^{+})^2}, \beta_{i,t}^{-} = \frac{\sum_{\tau=1}^{O} r_{i,\tau} r_{M,\tau}^{-}}{\sum_{\tau=1}^{O} (r_{M,\tau}^{-})^2}, \tag{4}$$

Different from the four realized semi-betas proposed by Bollerslev et al. (2021), the realized upside and downside betas condition only on the sign of market returns to account for joint asymmetric dependencies, which can be considered as a weighted sum of realized semi-betas,

$$\beta_{i,t}^{+} = \left(\beta_{i,t}^{P} - \beta_{i,t}^{M+}\right) \frac{\sum_{\tau=1}^{O} {r_{M,\tau}}^2}{\sum_{\tau=1}^{O} (r_{M,\tau}^{+})^2}, \beta_{i,t}^{-} = \left(\beta_{i,t}^{N} - \beta_{i,t}^{M-}\right) \frac{\sum_{\tau=1}^{O} {r_{M,\tau}}^2}{\sum_{\tau=1}^{O} (r_{M,\tau}^{-})^2}, \tag{5}$$

Note that the weights on the realized semi-betas only involve functions of market returns, which do not vary in the cross section.

Figure 3 plots the time series of realized betas, four realized semi-betas and realized upside and downside betas. We can see that different beta components show distinctly different patterns of variation over time. Take the four realized semi-betas as examples, during the 2008-09 financial crisis periods, though β^{M+} and β^{M-} fluctuated less pronouncedly than β^{P} and β^{N} in general, some sudden and substantial "spikes" in their levels occurred. Therefore, it would be interesting to examine whether these different beta components can provide additional predictive power for realized betas.

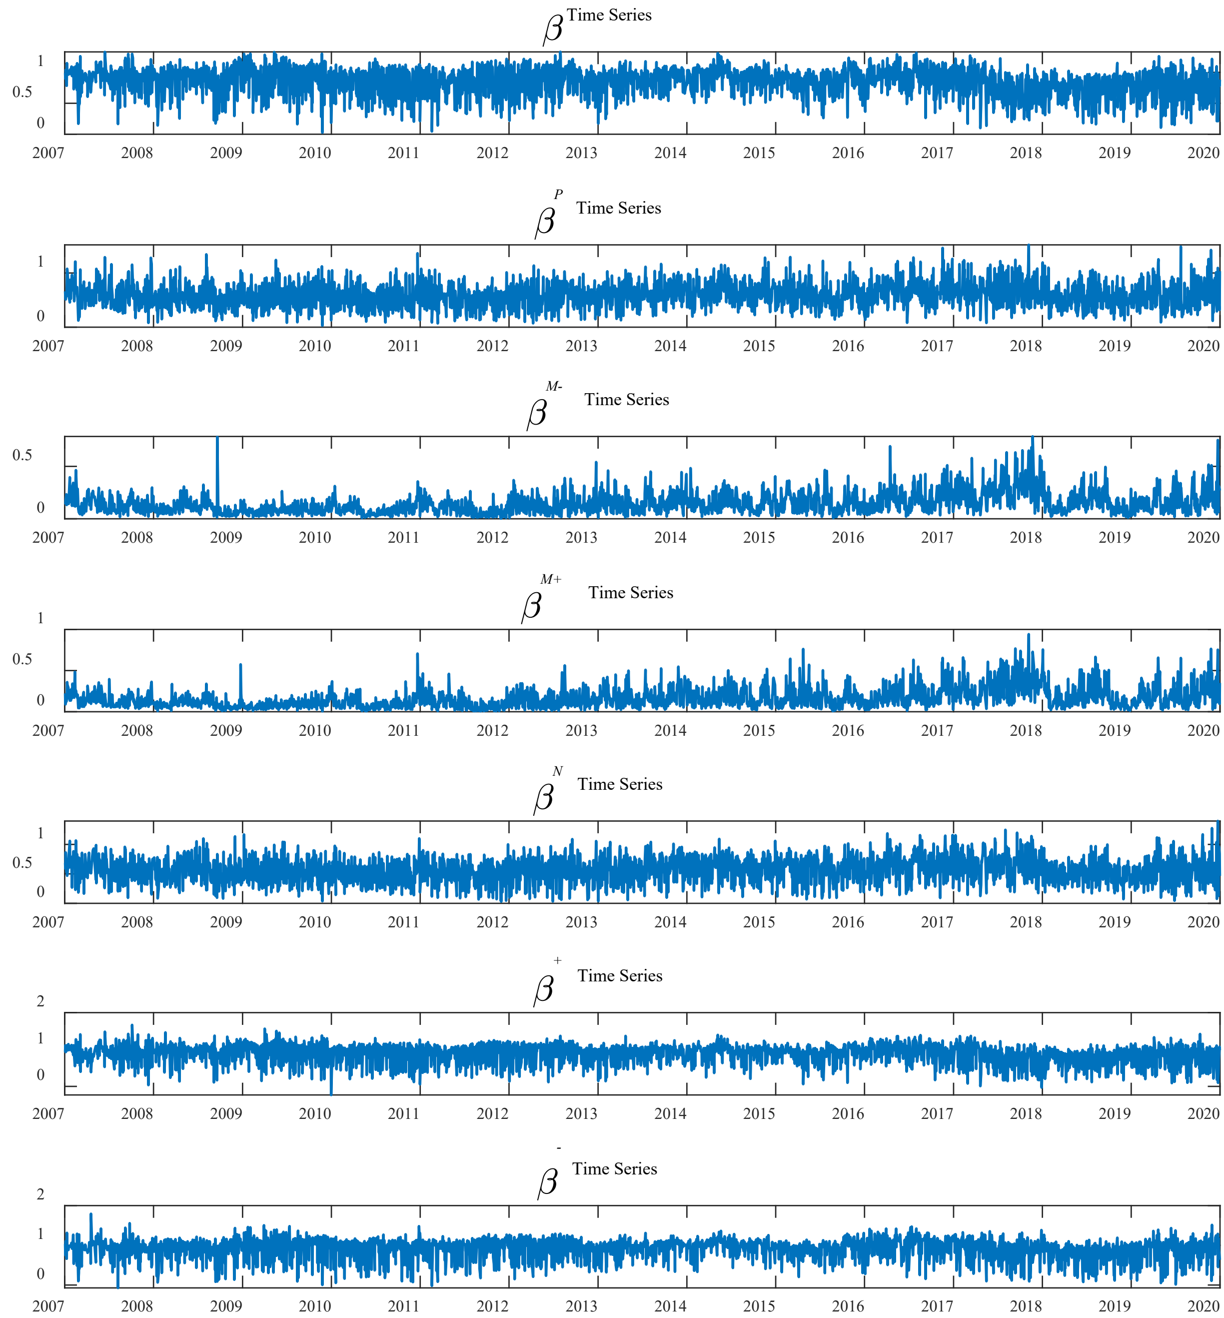

Figure 3. Time-series plots of the realized measures of beta and beta components. This figure plots the daily realized betas, four realized semi-betas, and realized upside and downside betas, averaged across all individual stocks.

The summary statistics of seven realized measures of beta and beta components are documented in Table 2. Their mean, median and standard deviation values are reported in Panel A of Table 2. It is confirmed that β^{M+} and β^{M-} are much less variable than others. In addition, Panel B of Table 2 shows the sample correlations. It is observed that the classic realized beta is strongly correlated with β^{N}, β^{+} and β^{-}, while the four realized semi-betas are much less correlated with the realized upside and downside betas contemporarily. Therefore, realized semi-betas and realized upside and downside betas may provide unique information to some extent on future realized betas.

Table 2

Summary statistics. Panel A reports the sample mean, median and standard deviation for each of the realized measures of beta and beta components. Panel B reports the sample correlations among them.

	β	β^p	β^{M-}	β^{M+}	β^N	β^+	β^-
Panel A: Summary statistics							
Mean	0.854	0.589	0.141	0.158	0.564	0.892	0.921
Median	0.898	0.572	0.116	0.124	0.569	0.959	0.986
St. Dev.	0.220	0.223	0.098	0.120	0.207	0.255	0.254
Panel B: Correlation							
β	1.000	0.492	0.078	0.031	0.585	0.601	0.562
β^p		1.000	0.306	0.681	-0.013	0.188	0.356
β^{M-}			1.000	0.597	0.573	0.153	-0.154
β^{M+}				1.000	0.163	-0.201	0.172
β^N					1.000	0.391	0.240
β^+						1.000	-0.236
β^-							1.000

Panel A of Figure 4 visualizes the empirical distributions of realized measures of beta and various beta components. The empirical distributions of realized betas and realized upside and downside betas are left-skewed with centers around unity. For realized semi-betas, we observe that β^{M+} and β^{M-} are symmetrically distributed, while β^P and β^N are significantly right-skewed. Panel B of Figure 4 plots the sample Autocorrelation Functions (ACFs) for all realized measures. Interestingly, β^{M+} and β^{M-} show the strongest long-memory features. The autocorrelations of other realized measures, including the realized beta, however, decline rapidly at first but soon climb to and stabilize at around between 0.05 and 0.2. The sample ACFs evidently show that the persistent structures of beta and its various components are considerably different from those of the realized volatility (see, e.g., Corsi, 2009).

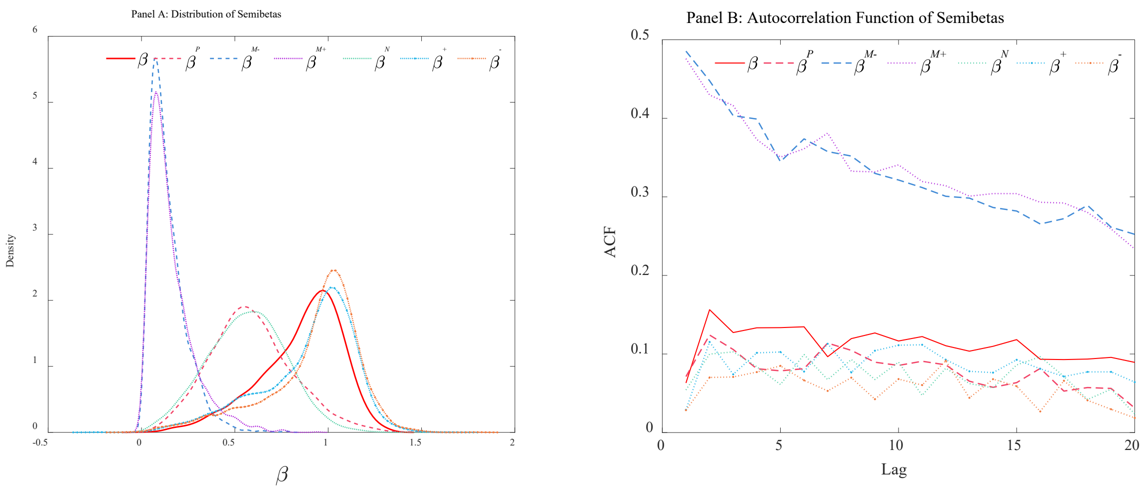

Figure 4. Empirical distributions and autocorrelation functions. Panel A plots the empirical distributions of realized measures of beta and various beta components. Panel B plots their sample autocorrelations.

2.4. Structural breaks

The aforementioned time-series analysis of realized betas, realized semi-betas, and realized upside and downside betas (e.g., Figure 3) provides us a further idea of improving the predictive accuracy by taking into account the structural breaks that may occur from time to time. We adopt Iterated Cummulative Sum of Square (ICSS) algorithm based on the Cummulative Sum (CUSUM) statistic proposed by Inclan and Tiao (1994) to detect multiple structural break points in the variance of daily realized measures of beta and beta components.

Figure 5 displays the testing results for structural breaks of realized betas, realized semi-betas and realized upside and downside betas (averaged across all individual stocks). It is clear that both concordant semi-beta components β^{P} and β^{N} have the lowest number (i.e., $K = 1$) of structural breaks. Other realized beta components β, β^{M-}, β^{M+}, β^{+} and β^{-} have three, five, ten, five and three structural breaks, respectively. Take the main research objective β as an example, the three detected structural break points correspond to the financial crisis of 2008, the European debt crisis of 2011, and the Fed interest rate hike of 2018, respectively. These structural breaks usually stem from important macroeconomic shocks and may in turn lead to significant future variations. For each of the individual stocks, one can reasonably presume that a much greater number of structural breaks would occur in betas. In short, taking the structural break dummies into consideration may provide additional predictive power for individual betas.

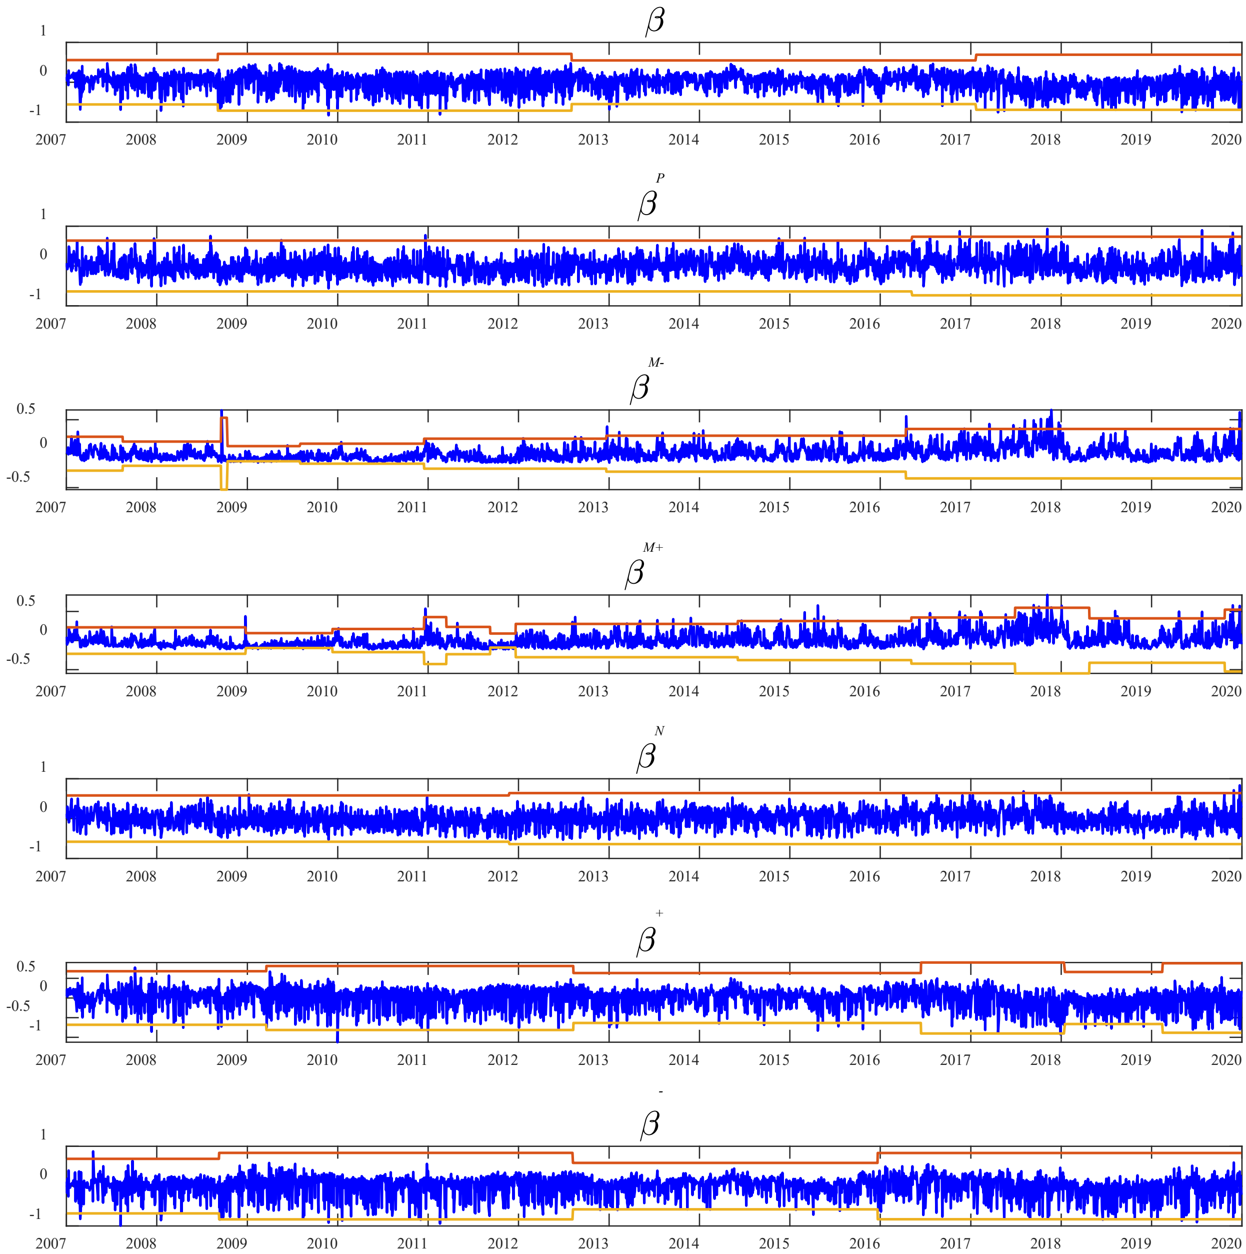

Figure 5. Structural break tests for realized betas, realized semi-betas, and realized upside and downside betas. This figure shows the number of structural breaks for each of the realized measures of beta and various beta components (averaged across all individual stocks) using the method of ICSS algorithm introduced by Inclan and Tiao (1994). The results show that β, β^{P}, β^{M-}, β^{M+}, β^{N}, β^{+} and β^{-} have three, one, five, ten, one, five and three structural breaks, respectively.

3. Forecasting Models

With the empirical properties of realized betas, realized semi-betas, and realized upside/downside betas shown in Section 2, we propose a set of new HAR-type forecasting models (which are denoted as HAR_{ARE}, HAR_{SB}, $HAR_{ARE.SB}$, and $LHAR_{ARE.SB}$, respectively) and compare them with three groups of competing models: 1) Difference-stationary models (including the random walk model); 2) Short-memory models (including AR(p), ARMA(p,q) and HAR models); 3) Long-memory models (including FI(d) and ARFIMA(p,d,q) models).

3.1. Random walk model

The simplest difference-stationary model is the Random Walk (RW) model without drift. In this case, the h-step-ahead beta forecasts are obtained by

$$\beta_{i,t+h} = \beta_{i,t+h-1} + \epsilon_{i,t+h}, \tag{6}$$

where $\epsilon_{i,t+h}$ is the mean-zero error term for stock i on day $t+h$. Note that we can rewrite Equation (6) as $\beta_{i,t+h} = \beta_{i,t} + \epsilon_{i,t+1} + \epsilon_{i,t+2} + \cdots + \epsilon_{i,t+h}$. Therefore, the (mean-squared) optimal h-step-ahead forecasts are directly given by the level of β_i on day t.

3.2. AR(p) model

The Autoregressive (AR) model is widely used for forecasting betas (see, e.g., Hooper et al., 2008, Reeves and Wu, 2013, and Cenesizoglu et al., 2017). For a p-order AR model, it allows us to predict the h-step-ahead $\beta_{i,t+h}$ for stock i based on previous p beta estimates of $\{\beta_{i,t}, \beta_{i,t-1}, \cdots, \beta_{i,t-p+1}\}$. Recall that since an AR(p) process has an exponentially decreasing autocorrelation function, it is referred to as a short-memory model (Hosking, 1996). The AR(p) model can be written as follows,

$$\beta_{i,t+h} = \alpha_i + \theta_{i,1}\beta_{i,t} + \ldots + \theta_{i,p}\beta_{i,t-p+1} + \epsilon_{i,t+h}, \tag{7}$$

where $\theta_{i,j}$ is the autoregressive coefficient of $\beta_{i,t-j+1} (j = 1,2, \cdots, p)$ and $\epsilon_{i,t+h}$ is the mean-zero error term for stock i on day $t+h$. The lag order p determined by AIC criterion is almost always one. Therefore, we simply set $p = 1$ and thus only focus on the AR(1) model in the forecasts.

3.3. ARMA(p,q) model

Another simple short-memory model that is widely used in the literature is the Autoregressive Moving Average (ARMA) model, which explicitly takes the exogenous disturbances $\{\varepsilon_{i,t}, \varepsilon_{i,t-1}, \cdots, \varepsilon_{i,t-q+1}\}$ into account. The ARMA(p,q) model is shown in Equation (8),

$$\beta_{i,t+h} = \alpha_i + \theta_{i,1}\beta_{i,t} + \ldots + \theta_{i,p}\beta_{i,t-p+1} + \delta_{i,1}\varepsilon_{i,t} + \ldots + \delta_{i,p}\varepsilon_{i,t-q+1} + \epsilon_{i,t+h}, \tag{8}$$

where $\theta_{i,j}$ is the autoregressive coefficient of $\beta_{i,t-j+1}$ $(j = 1,2, \cdots, p)$, $\delta_{i,k}$ is the moving average coefficient of $\varepsilon_{i,t-k+1}$ $(k = 1,2,\cdots,q)$, and $\epsilon_{i,t+h}$ is the mean-zero error term for stock i on day $t+h$. Both p and q are set to be ones according to preliminary AIC and BIC tests with a maximum lag length of $12[(T/100)^{0.25}]$. Hosking (1984) argued that the ARMA(1,1) model with roots lying near the unit circle can generate a time-series with long-memory properties. Even in such case, however, it still differs from the true long-memory model.

3.4. ARFIMA(p,d,q) and *FI(d) models*

The true long-memory models that are most used in the literature include Autoregressive Fractionally Integrated Moving Average (ARFIMA) model (Granger and Joyeux, 1980; Hosking, 1981) and Fractionally Integrated (FI) model. With a compact representation using lag polynomials, the parameter $\hat{d}_{i,T}$ in the differencing operator for ARIMA model must be an integer, while ARFIMA(p,d,q) model allows the differencing parameter to take fractional values, which is the determinant of long-term dependence. The long-memory properties of ARFIMA(p,d,q) model were empirically examined by Ding et al. (1996), Stock and Watson (2002) and many others. ARFIMA(p,d,q) model can be written as,

$$\gamma(L)(1-L)^{\hat{d}_{i,T}}\beta_{i,t} = \psi(L)\varepsilon_{i,t}, \tag{9}$$

where L is the lag operator, $(1-L)^{\hat{d}_{i,T}}$ is the fractional differencing operator, and $\varepsilon_{i,t}$ is the white noise disturbance. Within the fractional differencing operator, the parameter $\hat{d}_{i,T}$ denotes the order of differencing, which is estimated by the GPH method in this paper with a bandwidth of $\lambda = T^{0.5}$ for stock i over an estimation window with length T. $\gamma(L)$ and $\psi(L)$ are p-order and q-order stationary hysteretic polynomial operators, respectively, which can be expressed as follow,

$$\gamma(L) = 1 - \gamma_1 L - \gamma_2 L^2 - \cdots - \gamma_p L^p, \tag{10}$$

$$\psi(L) = 1 - \psi_1 L - \psi_2 L^2 - \cdots - \psi_q L^q. \tag{11}$$

The ARFIMA(p,d,q) model produces time series with hyperbolically decaying autocorrelations, while short-memory models, such as ARMA, can only produce exponentially or geometrically decaying autocorrelations (Bhardwaj and Swanson, 2006; Aye et al., 2014). Such a feature of ARFIMA can be easily verified by noting that,

$$(1-L)^{\hat{d}_{i,T}} = \sum_{j=0}^{t-1}(-1)^j \binom{\hat{d}_{i,T}}{j}(L)^j = 1 - \hat{d}_{i,T}L + \frac{\hat{d}_{i,T}(\hat{d}_{i,T}-1)}{2!}L^2 - \frac{\hat{d}_{i,T}(\hat{d}_{i,T}-1)(\hat{d}_{i,T}-2)}{3!}L^3 + \cdots + (-1)^{t-1}\frac{\hat{d}_{i,T}(\hat{d}_{i,T}-1)(\hat{d}_{i,T}-2)\cdots(\hat{d}_{i,T}-t+2)}{(t-1)!}L^{t-1} = \sum_{j=0}^{t-1} b_{i,j}(\hat{d}_{i,T}),\ with\ t = 1,2,\cdots,T, \tag{12}$$

for any $d > 1$. For $d < 0$, the difference filter can also be developed further using a hyper geometric function as follows,

$$(1-L)^{\hat{d}_{i,T}} = \Lambda(-\hat{d}_{i,T})\sum_{j=0}^{t-1} L^j\, \Lambda(j-\hat{d}_{i,T})/\Lambda(j+1) = F(-\hat{d}_{i,T}, 1, 1, L), \tag{13}$$

where $F(a,b,c,z) = \Lambda(c)/[\Lambda(a)\Lambda(b)]\sum_{j=0}^{t-1} z^j\, \Lambda(a+j)\Lambda(b+j)/[\Lambda(c+j)\Lambda(j+1)]$.

Combining Equation (9)-(13), one obtains a mean-zero ARMA(p,q) for the process of $\{\Delta^{\hat{d}_{i,T}}\beta_{i,t}\}$. Therefore, p and q can be determined by AIC or BIC criterion for ARMA models introduced in Section 3.3. The final prediction of $\beta_{i,t+h}$ can thus be obtained by using fractional inverse operation, $\beta_{i,t+h} = \Delta^{-\hat{d}_{i,T}}(\Delta^{\hat{d}_{i,T}}\beta_{i,t+h})$. Corresponding with AR(1) and ARMA(1,1) models, we also set $p = q = 1$ for ARFIMA model, i.e., we consider the ARFIMA(1,d,1) specification. FI(d) model is another widely-used long-memory model, which is a special case of ARFIMA(p,d,q) with $p = q = 0$.

3.5. HAR model

The Heterogeneous Autoregressive (HAR) model proposed by Corsi (2009) can also be used to characterize long-memory processes. For instance, many empirical studies on HAR models, including Andersen et al. (2007), Corsi et al. (2008), Andersen et al. (2011), etc., have successfully achieved the long-memory feature of realized volatility and thus yielded accurate out-of-sample volatility forecasts. Note that the classic HAR model is actually a short-memory model since it is essentially a constrained AR(22) model. The long-memory properties stem from the additive “cascade” structure of variable components defined over different time periods. Compared with ARFIMA and FI models, HAR model is parsimonious and much easier to implement. Recently, Becker et al. (2021) found that HAR model was a suitable choice for forecasting realized betas. For each realized beta series, HAR model can be expressed as follows,

$$\beta_{i,t+h} = a_i + \theta_{d,i}\beta_{i,t} + \theta_{w,i}\sum_{j=1}^{5}\frac{\beta_{i,t-j+1}}{5} + \theta_{m,i}\sum_{j=1}^{22}\frac{\beta_{i,t-j+1}}{22} + \epsilon_{i,t+h}, \tag{14}$$

where $\theta_{d,i}$, $\theta_{w,i}$ and $\theta_{m,i}$ are the coefficients of three predictors, i.e., daily realized beta, weekly realized beta, and monthly realized beta, respectively, for stock i on day t. Although HAR models are successful in modeling long-memory volatility processes, Baillie et al. (2019) showed that in general they can not capture the full scale of long-run dependence. It is thus possible to improve the classic HAR model by incorporating other predictors.

3.6. HAR$_{ARE}$ model

Based on the analysis in Section 2.3, realized semi-betas and realized upside and downside betas potentially capture some useful information regarding various asymmetric risk effects and thus may provide additional predictive power for realized betas. To account for these effects, we propose the HAR$_{ARE}$ model that replaces the predictors in classic HAR model (i.e., Equation (14)) with realized semi-betas and realized upside and downside betas. The HAR$_{ARE}$ model can be written as,

$$\beta_{i,t+h} = a_i + \theta_d{}'\beta_{t,d,i} + \theta_w{}'\beta_{t,w,i} + \theta_m{}'\beta_{t,m,i} + \epsilon_{i,t+h}, \tag{15}$$

$$\beta_{d,i,t} = \{\beta_{i,t}^{P} \quad \beta_{i,t}^{M-} \quad \beta_{i,t}^{M+} \; \beta_{i,t}^{N} \quad \beta_{i,t}^{+} \quad \beta_{i,t}^{-}\},$$

$$\beta_{w,i,t} = 1/5 \textstyle\sum_{j=1}^{5} \{\beta_{i,t-j+1}^{P} \quad \beta_{i,t-j+1}^{M-} \quad \beta_{i,t-j+1}^{M+} \; \beta_{i,t-j+1}^{N} \quad \beta_{i,t-j+1}^{+} \quad \beta_{i,t-j+1}^{-}\},$$

$$\beta_{m,i,t} = 1/22 \textstyle\sum_{j=1}^{22} \{\beta_{i,t-j+1}^{P} \quad \beta_{i,t-j+1}^{M-} \quad \beta_{i,t-j+1}^{M+} \; \beta_{i,t-j+1}^{N} \quad \beta_{i,t-j+1}^{+} \quad \beta_{i,t-j+1}^{-}\},$$

where $\beta_{d,i,t}$, $\beta_{w,i,t}$ and $\beta_{m,i,t}$ are defined, respectively, as the daily, weekly and monthly beta predictor vectors for stock i, and θ_d, θ_w and θ_m are their corresponding coefficient vectors.

3.7. HAR$_{SB}$ model

As discussed in Section 2.4, structural breaks are often observed in the time series of daily realized betas. To account for this prevailing feature, we propose another HAR variant, which is denoted as the HAR$_{SB}$ model. In particular, the occurrence of structural breaks is detected by ICSS algorithm within the in-sample estimation period. Specifically, the HAR$_{SB}$ model can be expressed as follows,

$$\beta_{i,t+h} = a_i + \theta_{d,i}\beta_{i,t} + \theta_{w,i} \textstyle\sum_{j=1}^{5} \frac{\beta_{i,t-j+1}}{5} + \theta_{m,i} \sum_{j=1}^{22} \frac{\beta_{i,t-j+1}}{22} + \sum_{n=1}^{N_i} \omega_{n,i} D_{n,i,t} + \epsilon_{i,t+h}, \tag{16}$$

where $D_{n,i,t}$ denotes a dummy variable that takes the value of one if a structural break is detected at time t for stock i and zero otherwise, and N_i counts the total number of structural breaks over the in-sample period for stock i.

3.8. *HAR$_{ARE.SB}$ model*

To take both asymmetric risk effects and structural breaks into consideration, we additionally propose the HAR$_{ARE.SB}$, which is naturally written as,

$$\beta_{i,t+h} = a_i + \theta_d{}'\beta_{t,d,i} + \theta_w{}'\beta_{t,w,i} + \theta_m{}'\beta_{t,m,i} + \sum_{n=1}^{N} \omega_{n,i} D_{n,i,t} + \epsilon_{i,t+h}, \tag{17}$$

$$\beta_{d,i,t} = \{\beta_{i,t}^{P} \quad \beta_{i,t}^{M-} \quad \beta_{i,t}^{M+} \; \beta_{i,t}^{N} \quad \beta_{i,t}^{+} \quad \beta_{i,t}^{-}\},$$

$$\beta_{w,i,t} = 1/5 \sum_{j=1}^{5} \{\beta_{i,t-j+1}^{P} \quad \beta_{i,t-j+1}^{M-} \quad \beta_{i,t-j+1}^{M+} \; \beta_{i,t-j+1}^{N} \quad \beta_{i,t-j+1}^{+} \quad \beta_{i,t-j+1}^{-}\},$$

$$\beta_{m,i,t} = 1/22 \sum_{j=1}^{22} \{\beta_{i,t-j+1}^{P} \quad \beta_{i,t-j+1}^{M-} \quad \beta_{i,t-j+1}^{M+} \; \beta_{i,t-j+1}^{N} \quad \beta_{i,t-j+1}^{+} \quad \beta_{i,t-j+1}^{-}\}.$$

where $\beta_{d,i,t}$, $\beta_{w,i,t}$ and $\beta_{m,i,t}$ are defined, respectively, as the daily, weekly and monthly beta predictor vectors for stock i, and θ_d, θ_w and θ_m are their corresponding coefficient vectors. All the parameters in Equation (17) are in line with those of HAR$_{ARE}$ model and HAR$_{SB}$ model. Although the HAR$_{ARE.SB}$ model takes into account both internal (i.e., asymmetric risk effects) and external (i.e., structural breaks) factors to improve the forecasting accuracy, it raises the potential issue of overfitting. In addition, the model structure of realized betas can be time-varying and can differ across individual assets. Therefore, we introduce the LASSO method to solve these problems in next section.

3.9. *LASSO approach* and *LHAR$_{ARE.SB}$ model*

One of the main potential problems of the three proposed HAR variants, i.e., HAR$_{ARE}$, HAR$_{SB}$ and HAR$_{ARE.SB}$ models, is that the total number of model parameters is relatively large, which may lead to the overfitting issue. For the sake of model parsimony, we employ the Least Absolute Shrinkage and Selection Operator (LASSO) method proposed by Tibshirani (1996) for variable selection in this paper. LASSO utilizes a penalty that operates as a function of regression coefficients, so that it effectively controls the model complexity and avoids overfitting. Rather than minimizing the loss function associated with an ℓ_2-norm penalty on the coefficients (such as the ridge regression), LASSO employs an ℓ_1-norm penalty that sets an upper bound on the sum of absolute values of all the coefficients, i.e., $P(\boldsymbol{\theta}) = \lambda \sum_{i=1}^{P} |\theta_i|$. Therefore, LASSO minimizes the following loss function,

$$L(\theta_1, \theta_2, \cdots, \theta_p) = \left\| \mathbf{Y} - \sum_{i=1}^{P} \mathbf{X}_i \theta_i \right\|^2 + \lambda \sum_{i=1}^{P} |\theta_i|. \tag{18}$$

For any $\lambda > 0$, the ℓ_1-norm penalty generates the desired sparsity, i.e., some regression coefficients are shrunken to exactly zeros. In other words, predictors that lack of sufficient power are eliminated. Particularly, the λ for the HAR model is selected by a 10-fold cross-validation.

Therefore, the model we propose in this paper is a combination of the LASSO approach and the $\text{HAR}_{\text{ARE.SB}}$ model, which is denoted as $\text{LHAR}_{\text{ARE.SB}}$. By using a rolling-window estimation scheme, the selected predictors for $\text{LHAR}_{\text{ARE.SB}}$ model would be different over time. It may also be interesting to study how the importance of asymmetric risk effects and structural breaks for forecasting betas.

4. Forecasting Results

To compare the predictive performance of different models and, in particular, to examine the importance of structural breaks and asymmetric risk effects for forecasting realized betas, we carry out an extensive forecasting "horse-race" using all the models introduced in Section 3. In our out-of-sample forecasts, we set the rolling windows of 600 days, which is approximately 1/5 of the full-sample length. We also conduct the robust analysis by specifying the rolling window of 800 days and 1000 days, which correspond to 1/4 and 1/3 of the full-sample length. We present the short-run forecast results (h = 1) in this section and also conduct robust analysis regarding mid-term (h = 5) and long-term forecasts(h = 22). The predictive performance of all competing models is evaluated using both statistical and economic measures. Predictive accuracy tests are also conducted and the testing results are reported in the analysis.

4.1. Statistical and economic criteria

We consider two different statistical loss functions to compare the predictive accuracy of various models (Andersen et al., 1999), including the Root Mean Squared Error (RMSE) and the Mean Absolute Error (MAE) as shown below,

$$\text{RMSE} = \sqrt{\frac{1}{\vartheta} \sum_{t=1}^{\vartheta} \left(\beta_{i,t+h} - \hat{\beta}_{i,t+h}\right)^2}, \tag{19}$$

$$\text{MAE} = \frac{1}{\vartheta} \sum_{t=1}^{\vartheta} \left|\beta_{i,t+h} - \hat{\beta}_{i,t+h}\right|, \tag{20}$$

where ϑ is the total number of beta forecasts over the out-of-sample period, $\beta_{i,t+h}$ is the true value of realized beta on day t+h and $\hat{\beta}_{i,t+h}$ is the corresponding ex ante forecast.

Besides reporting the four statistical measures, we conduct pairwise Diebold–Mariano (DM) predictive accuracy tests using all of the aforementioned loss functions. DM tests assume a null hypothesis of equal predictive ability between a pair of competing models and are asymptotically normally distributed. Given a certain significance level, e.g., $\alpha = 5\%$, if the null hypothesis is rejected, we conclude that the two competing models have a different forecasting performance. The sign of DM test statistics can be further used to determine which of two competing models is statistically significantly better. In addition to pairwise model comparison, we also carry out the Model Confidence Set (MCS) test proposed by Hansen et al. (2011) for multiple comparisons. One of the main appealing features of MCS procedure when compared to other predictive accuracy tests is that researchers do not need to specify a benchmark model. With a certain level of confidence, e.g., $(1-\alpha) = 95\%$, the MCS procedure eliminates inferior models sequentially and formulates a final set of best models. To calculate the MCS p-values, we use the block bootstrap method with 10000 replications. The greater the p value of a competing model achieves, the higher the probability of being included in the final confidence set of best models.

To address the potential errors-in-variables issue in the forecast evaluation proxy, we construct portfolios similar to Fama and MacBeth (1973) and Hollstein and Prokopczuk (2016). Specifically, we sort all individual stocks in ascending order based on their realized beta values obtained during the second-to-last non-overlapping beta estimation window. In doing so, we follow Becker et al. (2021) to use a common sorting variable for all forecasting models.[4] In this way, all the individual stocks are sorted into 20 portfolios.[5] To yield the portfolio beta forecast β_p, we weight each individual beta forecast based on the firm's market value at the time the forecast is generated, i.e., $\beta_p = \sum_{i=1}^{N_p} w_{i,p}\,\beta_i$, where $w_{i,p}$ is the market capitalization and N_p is the number of stocks in portfolio p. The data of stocks' market capitalization are obtained from JoinQuant database. Figure

[4] Similar to Becker et al. (2021), we do so to ensure that (i) there is a spread in the market betas of the different portfolios. (ii) Using a sorting variable that is independent of the predictor variables is important to avoid discriminating against any of the predictors. Otherwise, the stocks with the highest positive and negative noise for the estimator upon which the betas are sorted are likely to end up in the extreme portfolios and that noise cannot be fully diversified. For stocks with missing estimates for the sorting variable, we set it to unity.

[5] Let N be the total number of individual stocks and int(N/20) be the largest integer less than or equal to N/20. Each of the 18 portfolios in the middle has int(N/20) stocks. If N is even, each of the two portfolios at both ends has int(N/20)+1/2[N-int(N/20)] stocks; If N is odd, the last portfolio has an additional stock.

6 shows the average market capitalizations of the stocks in different portfolios.

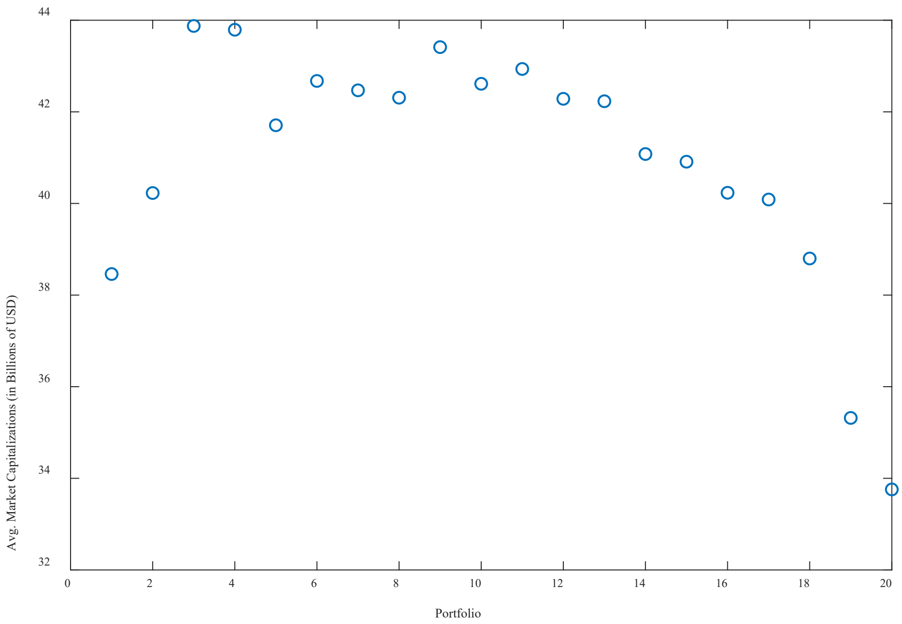

Figure 6. Average market capitalization of the stocks in each portfolio. This figure shows the average market capitalizations (in Billions of USD) of the stocks (sorted in ascending order) in the 20 portfolios. The first and the last portfolios have 19 and 20 stocks, respectively. On average, each of the rest portfolios contains 16 stocks.

It is also of great importance to assess beta forecasts based on some economic criteria since accurate beta forecasts are critical to asset allocation. First, similar to Becker et al. (2021), we evaluate different forecasting models based on their abilities to create portfolios arranged in strict order of ex-post realized betas. For each model, we first sort all individual stocks into ten portfolios based on their respective beta forecasts at the end of each day. Next we calculate the ex-post realized beta of each beta-sorted portfolio. A good forecasting model should be able to generate a monotonically increasing pattern in portfolios' realized betas. Moreover, the spread between the ex-post realized betas of high and low beta-sorted portfolios should be large.

Second, following Hollstein et al. (2020), we evaluate the ability of each model for creating market-neutral portfolios, by checking their average risk exposure and ex-post realized betas. To create such portfolios, we set the weight $v_{j,t}$ so that the equation $v_{j,t}\beta_{j,t}^{long} - \beta_{j,t}^{short} = 0$ is fulfilled, where $\beta_{j,t}^{long}$ and $\beta_{j,t}^{short}$ represent the beta forecast of long portfolio (consisting of top quantile of beta-sorted stocks) and the beta forecast of short portfolio (consisting of bottom quantile of beta-sorted stocks), respectively. Then, we compute the average risk exposure $\sum_{t=1}^{T}\frac{1}{T}\left|v_{j,t}\beta_{j,t}^{long}-\beta_{j,t}^{short}\right|$ for each model and test whether the ex-post realized beta of the anomaly portfolios deviates from zero on average.

4.2. Statistical evaluations

A number of clear-cut findings emerge upon inspection of the results contained in Tables 3-4. Table 3 summarizes the performance of various models for forecasting realized betas of the 20 constructed portfolios, based on the RMSE loss. Testing results of pairwise model comparison, as well as multiple comparisons are also reported.

First, according to the average RMSE values and the DM tests at the 5% level, the HAR-type models perform significantly better than the benchmark AR and RW model and slightly better than the true long-memory models (FI and ARFIMA). The results suggest that the short-memory and stationary models are not suitable for forecasting realized betas. In addition, the HAR-type models perform slightly better than the true long-memory models in general. The $\text{LHAR}_{\text{ARE.SB}}$ model performs the best and yield significantly higher precision than other models.

Second, the structural break dummies and decomposed beta components provide additional predictive power. For instance, HAR_{ARE}, HAR_{SB} and $\text{HAR}_{\text{ARE.SB}}$ generate lower RMSE values on average when compared to the classic HAR model. Turning to the DM test results, this finding is more distinct. For all of the 20 portfolios, HAR_{ARE}, HAR_{SB} and $\text{HAR}_{\text{ARE.SB}}$ significantly outperform the classic HAR based on DM tests at the 5% level. Fourth, when the LASSO approach is conducted on $\text{HAR}_{\text{ARE.SB}}$, a further significant improvement of predictive accuracy is observed. For 14 out of 20 portfolios, $\text{LHAR}_{\text{ARE.SB}}$ achieves the lowest RMSEs from amongst all models.

Finally, the MCS testing results also support the superior performance of $\text{LHAR}_{\text{ARE.SB}}$ as it is always included in the confidence set of “best” models with the level of either 90% or 80%. Interestingly, the long-memory models ARFIMA and FI are the only other alternatives that can be selected by MCS_{90} (MCS_{80}) tests. Table 4 demonstrates similar findings based on the MAE loss. For instance, ARFIMA, HAR_{ARE}, HAR_{SB} and $\text{HAR}_{\text{ARE.SB}}$ general outperform the baseline HAR model as well as the short-memory and stationary-difference models, and LASSO approach significantly improves the performance of the HAR-type models. However, under the MAE loss, ARFIMA shows relatively comparable performance when compared to the HAR variants, such as

HAR_{ARE}, HAR_{SB} and $HAR_{ARE.SB}$.

According to the results summarized in Tables 3-4, we conclude that it is of great significance to characterize the main empirical features of realized betas, such as long-run persistence, structural breaks and asymmetric risk effects, in a tractable and parsimonious way to yield accurate forecasts.

Table 3

Out-of-sample forecasting results based on RMSE loss. This table shows the results of out-of-sample short-term (h = 1) prediction based on RMSE loss function with rolling-window scheme. The first row is the average RMSEs of all 20 portfolios. The next row "Best" refers to the frequencies at which lowest RMSEs are achieved for each model when forecasting portfolio betas. MCS_{90} (MCS_{80}) shows the number of times a forecasting model is included within the 90% (80%) confidence set under the RMSE loss. Finally, the rows denoted by "vs. X" demonstrate the pairwise DM test results. For instance, the integer "13" (in Row 5, Column 2) indicates that for 13 out of 20 portfolios, the AR model significantly beats the random walk model based on the DM test at the 5% level of significance.

	RW	AR	ARMA	ARFIMA	FI	HAR	HAR_{ARE}	HAR_{SB}	$HAR_{ARE.SB}$	$LHAR_{ARE.SB}$
RMSE	0.4236	0.4586	0.4160	0.2856	0.3069	0.2454	0.2332	0.2305	0.2247	0.1797
Best	0	0	0	6	0	0	0	0	0	14
MCS_{80}	0	0	0	7	1	0	0	0	0	16
MCS_{90}	0	0	0	7	1	0	0	0	0	17
vs. RW	0	13	15	17	17	20	20	20	20	20
vs. AR	7	0	20	20	20	19	19	19	19	19
vs. ARMA	5	0	0	20	20	16	16	16	16	19
vs. ARFIMA	3	0	0	0	6	6	9	9	9	12
vs. FI	3	0	0	14	0	8	10	10	11	14
vs. HAR	0	1	4	14	12	0	20	20	20	20
vs. HAR_{ARE}	0	1	4	11	10	0	0	19	20	20
vs. HAR_{SB}	0	1	4	11	10	0	1	0	20	20
vs. $HAR_{ARE.SB}$	0	1	4	11	9	0	0	0	0	20
vs. $LHAR_{ARE.SB}$	0	1	1	8	6	0	0	0	0	0

Table 4

Out-of-sample forecasting results based on MAE loss. This table shows the results of out-of-sample short-term (h = 1) prediction based on MAE loss function with rolling-window scheme. See the notes to Tables 3 for further details.

	RW	AR	ARMA	ARFIMA	FI	HAR	HAR_{ARE}	HAR_{SB}	$HAR_{ARE.SB}$	$LHAR_{ARE.SB}$
MAE	0.3205	0.3632	0.3369	0.2267	0.2500	0.1892	0.1784	0.1761	0.1709	0.1288
Best	0	0	0	3	0	0	0	0	0	17
MCS_{80}	0	0	0	5	0	0	0	0	0	18
MCS_{90}	0	0	0	6	0	0	0	0	0	18
vs. RW	0	11	15	17	16	20	20	20	20	20
vs. AR	9	0	20	20	20	20	20	20	20	20
vs. ARMA	5	0	0	20	20	17	19	19	19	20

vs. ARFIMA	3	0	0	0	4	9	9	9	9	14
vs. FI	4	0	0	16	0	9	12	12	12	20
vs. HAR	0	0	3	11	11	0	20	20	20	20
vs. $\mathrm{HAR_{ARE}}$	0	0	1	11	8	0	0	19	20	20
vs. $\mathrm{HAR_{SB}}$	0	0	1	11	8	0	1	0	20	20
vs. $\mathrm{HAR_{ARE.SB}}$	0	0	1	11	8	0	0	0	0	20
vs. $\mathrm{LHAR_{ARE.SB}}$	0	0	0	6	0	0	0	0	0	0

4.3. Economic analysis

In addition to the statistical evaluations presented thus far, we also compare various forecasting models based on the economic criterion introduced in Section 4.1. Table 5 shows the ex-post realized betas of ten beta-sorted portfolios constructed using ex ante beta forecasts yielded from each model. Although the monotonically increasing pattern in portfolios' betas is universally observed, the high-minus-low portfolio spreads can be very different across the ten competing models. For instance, as shown in the last row of Table 5, the entries corresponding to true long-memory models exceed 1.37. In particular, the ARFIMA model generates a large spread of 1.4466, while the short-memory and stationary-difference models can only generate spreads below 1.2. In addition, the HAR model as well as all of its variants generate relatively large spreads with values greater than 1.37. Amongst all the HAR-type models, $\mathrm{LHAR_{ARE.SB}}$ is the best-performing one that generates the largest spread of 1.4735.

We also evaluate the ability of each model for creating market-neutral portfolios. After sorting stocks into 10 portfolios based on the forecasted betas, we construct ex-ante beta-neutral portfolios by solving the equation $v_{j,t}\hat{\beta}_{j,t}^{long} - \hat{\beta}_{j,t}^{short} = 0$, where $\hat{\beta}_{j,t}^{short}$ is the beta of long portfolios and $\hat{\beta}_{j,t}^{short}$ is the beta of short portfolio. And applying the resulting weight $v_{j,t}$ constructs ex-ante beta neutral portfolios with $v_{j,t}\beta_{j,t}^{long} - \beta_{j,t}^{short}$. Following Papageorgiou et al.(2016) and Hollstein (2019), we present p-value of Newey West t-statistics to test whether ex post portfolio betas deviate from 0. The results are shown in Table 6.

The performance varies substantially across different forecasting models. The portfolios constructed using short-memory models in general have relatively low risk exposures. For instance, ARMA achieves an average risk exposure of 0.2076, which is almost identical to the performance of $\mathrm{LHAR_{ARE.SB}}$ model. For the true long-memory models, FI also shows a good performance, ARFIMA, however, constructs portfolios with large risk exposures on average. The HAR model and our proposed HAR-type models show similar performance when compared to the FI model. Amongst all competing models, $LHAR_{ARE.SB}$ works best in terms of this economic criterion. Next, turn to the *t*-test results for the null of zero ex-post portfolios' beta. We observe that only the HAR-type models can generate *p*-values greater than 0.1. The testing results show that neither ARMA model nor FI model constructs portfolios with zero ex-post beta. So we can conclude that $LHAR_{ARE.SB}$ constructs portfolios with lowest average risk exposures and these portfolios are market-neutral on average.

Table 5

Ex-post realized betas of portfolios sorted by ex ante beta forecasts. This table presents the ex-post realized betas of portfolios sorted by ex ante beta forecasts generated from each model. At the end of each trading day, we sort the individual stocks into ten portfolios based on stocks' beta forecasts. In addition, we show the high-minus-low spreads, i.e., ex-post realized beta of the "highest-beta" portfolio minus that of the "lowest-beta" portfolio. The statistic values in parentheses are the HAC-robust standard errors (Andrews, 1991). *, ** and *** indicate significance at 10%, 5% and 1% levels, respectively.

	RW	AR	ARMA	ARFIMA	FI	HAR	HAR_{ARE}	HAR_{SB}	$HAR_{ARE.SB}$	$LHAR_{ARE.SB}$
1	0.2221***	0.4111***	0.2757***	0.0694***	0.0995***	0.0943***	0.0864***	0.0827***	0.0798***	0.0512***
	(0.0155)	(0.0135)	(0.0110)	(0.0142)	(0.0150)	(0.0156)	(0.0156)	(0.0157)	(0.0157)	(0.0161)
2	0.5146***	0.6218***	0.5942***	0.4549***	0.4841***	0.4872***	0.4837***	0.4814***	0.4768***	0.4649***
	(0.0089)	(0.0102)	(0.0114)	(0.0098)	(0.0094)	(0.0093)	(0.0093)	(0.0094)	(0.0094)	(0.0093)
3	0.6023***	0.6622***	0.6614***	0.5730***	0.5690***	0.5802***	0.5741***	0.5748***	0.5725***	0.5660***
	(0.0090)	(0.0111)	(0.0106)	(0.0092)	(0.0091)	(0.0088)	(0.0087)	(0.0087)	(0.0087)	(0.0088)
4	0.6720***	0.7184***	0.7249***	0.6689***	0.6677***	0.6646***	0.6642***	0.6636***	0.6622***	0.6588***
	(0.0092)	(0.0107)	(0.0116)	(0.0090)	(0.0088)	(0.0091)	(0.0091)	(0.0090)	(0.0090)	(0.0086)
5	0.7520***	0.7530***	0.7689***	0.7709***	0.7558***	0.7556***	0.7551***	0.7521***	0.7531***	0.7506***
	(0.0098)	(0.0105)	(0.0111)	(0.0095)	(0.0097)	(0.0091)	(0.0090)	(0.0090)	(0.0090)	(0.0089)
6	0.8219***	0.7816***	0.8250***	0.8785***	0.8528***	0.8439***	0.8485***	0.8489***	0.8495***	0.8528***
	(0.0099)	(0.0116)	(0.0118)	(0.0103)	(0.0103)	(0.0095)	(0.0096)	(0.0095)	(0.0094)	(0.0096)
7	0.9078***	0.8678***	0.8759***	0.9882***	0.9557***	0.9561***	0.9589***	0.9565***	0.9581***	0.9685***
	(0.0112)	(0.0117)	(0.0127)	(0.0114)	(0.0112)	(0.0110)	(0.0110)	(0.0108)	(0.0109)	(0.0108)
8	1.0314***	0.9032***	0.9618***	1.1379***	1.1033***	1.1029***	1.1071***	1.1096***	1.1120***	1.1285***
	(0.0132)	(0.0122)	(0.0128)	(0.0137)	(0.0135)	(0.0136)	(0.0136)	(0.0136)	(0.0136)	(0.0135)
9	1.2754***	1.0183***	1.1387***	1.4515***	1.4069***	1.4016***	1.4127***	1.4130***	1.4180***	1.4550***
	(0.0185)	(0.0155)	(0.0156)	(0.0196)	(0.0196)	(0.0187)	(0.0189)	(0.0189)	(0.0189)	(0.0192)
10	1.3246***	1.0323***	1.1634***	1.5160***	1.4696***	1.4697***	1.4806***	1.4819***	1.4866***	1.5247***
	(0.0195)	(0.0156)	(0.0160)	(0.0212)	(0.0212)	(0.0208)	(0.0210)	(0.0210)	(0.0210)	(0.0212)
10-1	1.1026***	0.6212***	0.8877***	1.4466***	1.3701***	1.3754***	1.3942***	1.3992***	1.4068***	1.4735***
	(0.0275)	(0.0195)	(0.0194)	(0.0298)	(0.0299)	(0.0300)	(0.0303)	(0.0302)	(0.0303)	(0.0309)

Table 6

Market-neutral anomaly portfolios. This table presents the average risk exposure of market-neutral anomaly portfolios and the testing results for the null of zero ex-post realized betas of these portfolios. The first row contains the values calculated by $\sum_{t=1}^{T}\frac{1}{T}\left|v_{j,t}\beta_{j,t}^{long}-\beta_{j,t}^{short}\right|$. The next row shows the p-values of one-sample t-tests for the null of zero ex-post realized beta of the anomaly portfolios. $*$, $**$, and $***$ indicate significance at the 10%, 5%, and 1% levels, respectively.

	RW	AR	ARMA	ARFIMA	FI	HAR	HAR_{ARE}	HAR_{SB}	$\text{HAR}_{\text{ARE.SB}}$	$\text{LHAR}_{\text{ARE.SB}}$
Risk exposure	0.4052***	0.2836***	0.2076***	0.3604***	0.2417***	0.2427	0.2370	0.2348	0.2320	0.2070
p-value	0.0000	0.0000	0.0000	0.0000	0.0000	0.4240	0.5944	0.6177	0.3373	0.1414

Market participants usually care more about the economic gains than forecast precision. Therefore, we design a profitable trading strategy based on beta models to identify undervalued (overvalued) stocks, as suggested by Bollerslev et al. (2023).

The portfolio construction procedure is as follows: First, we rely on the traditional CAPM model $\hat{r}_{i,t+1}=\left(r_{M,t}-r_{f,t}\right)\hat{\beta}_{i,t+1}+r_{f,t}$ to obtain the forecast return $\hat{r}_{i,t+1}$ of stock i for the next day, where $r_{M,t}$ and $r_{f,t}$ are the returns of the S&P 500 and the risk-free interest rate on the current day, respectively, and $\hat{\beta}_{i,t+1}$ is the beta forecast for the next day. We use the interest rate on the U.S. 10-year bond as the risk-free interest rate. Second, each day we select 10 stocks with the highest $\left|\hat{r}_{i,t+1}-r_{i,t}\right|$ and construct the long-short portfolio as follows: if the return forecast of stock i is above (below) the actual current-day return, we expect its price to rise (fall) the next day and take a long (short) position, allocating w_i amount of money, where w_i is the weight based on the market capitalization of stock i. Finally, we repeat this strategy daily, using the daily beta forecasts to compute the cumulative returns, average annualized returns, and Sharpe ratio to assess the performance of different forecasting models. The intuition behind this trading strategy is straightforward: rational investors aim to buy undervalued stocks and sell overvalued stocks. The differences between current returns and return forecasts indicate arbitrage opportunities.

To further evaluate the economic significance of the forecasting models, we employ the quadratic utility to compute the economic value $\triangle$, according to Bollerslev et al. (2016b):

$$U(r^k,\gamma)=\left(1+r_p^k\right)-\frac{\gamma}{2(1+\gamma)}(1+r^k)^2, \tag{21}$$

where r_p^k refers to the monthly cumulative rate of return on portfolio associated with model k, and γ refers to the risk aversion rate. In this paper, we consider two levels of risk aversion: a mild rate $\gamma = 1$ and a strong rate $\gamma = 10$. The economic value Δ is the value such that

$$\sum_{t=T_1+1}^{T} U(r^k, \gamma) = \sum_{t=T_1+1}^{T} U(r^l - \Delta, \gamma). \tag{22}$$

The greater the Δ, the more returns a risk-aver investor is willing to sacrifice to switch from model l to model k. For brevity, we display only the economic values of each forecasting model against the benchmark RW model.

Table 7 presents the results of the economic evaluation for each forecasting model based on our stock selection strategy. The results show that the portfolio based on the $\text{LHAR}_{\text{ARE.SB}}$ model achieves the highest annualized return, Sharp ratio, and economic values compared to those based on other models, with the average annualized return of 20.0659% and the average Sharp ratio of 0.2049. Notably, almost all long-memory models perform better than short-memory models and the benchmark models. Furthermore, according to the results of $\triangle$ for $\gamma = 1$ and $\gamma = 10$, the economic gains of the $\text{LHAR}_{\text{ARE.SB}}$ model are more pronounced when the risk aversion rate is higher. This suggests that the utility for risk-seeking investors can be significantly improved by relying on the $\text{LHAR}_{\text{ARE.SB}}$ model for stock trading.

Figure 7 plots the cumulative returns of the proposed HAR-type models and the S&P 500 over the out-of-sample period. We find that the $\text{LHAR}_{\text{ARE.SB}}$ model yields better out-of-sample cumulative returns compared to other models. Additionally, out-of-sample cumulative portfolio returns based on the HAR-type models outperform the benchmark S&P 500 index return.

Table 7

Economic values evaluation. This table presents the annualized cumulative returns, average returns, Sharp ratios, and economic values $\triangle$ for the forecasting models. γ refers to the risk aversion rate. We use the interest rate on U.S. 10-year bonds as the risk-free interest rate. The economic value $\triangle$ is estimated using the quadratic utility with risk aversion, as described by Bollerslev et al. (2016b).

Models	Annualized Return (%)		Sharp Ratio		$\triangle$	
	Cumulative	Average	Cumulative	Average	γ=1	γ=10
RW	-0.5899	-0.4341	-0.0276	-0.0259	-	-
AR	-12.0801	-11.9234	-0.1567	-0.1550	-11.5682	-12.2773
ARMA	-13.8240	-13.6656	-0.1754	-0.1736	-13.4808	-15.8151
ARFIMA	3.7465	3.9091	0.0208	0.0226	3.6673	-3.5047
FI	-7.8423	-7.6834	-0.1083	-7.6834	-7.5522	-10.4267
HAR	15.1556	15.3108	0.1501	0.1519	15.7947	16.2305
HAR_{ARE}	17.6299	17.7862	0.1774	0.1792	18.1585	17.6047
HAR_{SB}	15.2476	15.4063	0.1495	0.1513	15.5459	12.7565

HAR$_{ARE.SB}$	14.4275	14.5845	0.1411	0.1429	14.8946	13.7678
LHAR$_{ARE.SB}$	19.9098	20.0659	0.2032	0.2049	20.4597	20.1002

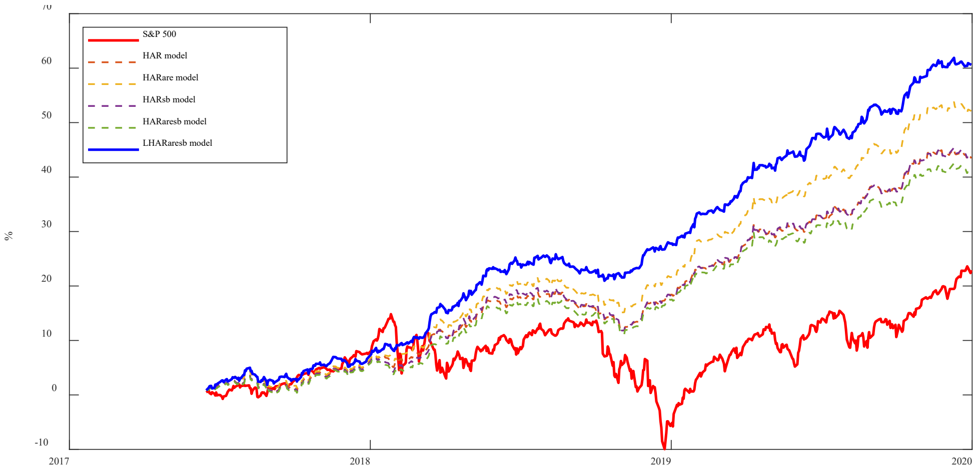

Figure 7. Accumulative returns of various models and S&P 500 for the out-of-sample period. This figure shows the accumulative returns for various HAR-type models as well as for S&P 500 index for the out-of-sample period.

5. Further Analysis

5.1 Decomposition of MSE

In this section, we delve into the reasons behind the wide differences in forecasting performance of various models. Following Mincer and Zarnowitz (1969), the analysis approach is based on decomposing the Mean Squared Error (note that MSE = RMSE^2) of out-of-sample realized beta forecasts as follows,

$$MSE_i = (\bar{\beta}_i - \bar{\hat{\beta}}_i)^2 + (1 - b_i)^2 \sigma^2(\hat{\beta}_i) + (1 - p_i^2)\sigma^2(\hat{\beta}_i), \tag{23}$$

where $\bar{\hat{\beta}}_i$ and $\sigma^2(\hat{\beta}_i)$ are the sample mean and variance of beta forecasts, respectively, and b_i and p_i^2 are the slope coefficient and the coefficient of determination of the regression $\beta_i = a_i + b_i\hat{\beta}_i + \epsilon_i$, respectively. The first component of MSE in Equation (23) is associated with the bias term, which indicates, on average, how much the forecasts deviate from the true values. The second term characterizes inefficiency. A large inefficiency represents that the forecasting model tends to yield positive forecast errors for low values and negative forecast errors for high values. The last random error term is unrelated to the forecasts as well as the true values.

Table 8 reports mean values for each component of MSE. Recall from Table 3, RW, AR and ARMA models generate large RMSEs. However, their causes are very distinct. For instance, short-memory models, i.e., AR and ARMA, have sizable biases and random errors. On the contrary, RW model has surprisingly the smallest bias term, in the meantime, however, it has the largest inefficiency component amongst all models. Next, comparing ARFIMA with ARMA, we find that the true long-memory structure plays a key role in substantially reducing the bias and random error terms. Again, from Table 3, we have learned that HAR-type models perform exceedingly well in terms of RMSE. Based on the analysis of MSE decomposition, it is mainly because they are all approximately unbiased. The LASSO approach further reduces the inefficiency and random error components when applied to $\mathrm{HAR}_{\mathrm{ARE.SB}}$, though at a slight cost to bias term. As a consequence, it is not surprising to find that $\mathrm{LHAR}_{\mathrm{ARE.SB}}$ is the best-performing model in terms of RMSE.

Table 8

Decomposition of MSE. This table shows the results of MSE decomposition based on short-term (h = 1) forecast errors of realized betas. All entries within this table represent the average across 20 beta-sorted portfolios.

	RW	AR	ARMA	ARFIMA	FI	HAR	$\mathrm{HAR}_{\mathrm{ARE}}$	$\mathrm{HAR}_{\mathrm{SB}}$	$\mathrm{HAR}_{\mathrm{ARE.SB}}$	$\mathrm{LHAR}_{\mathrm{ARE.SB}}$
Bias	0.0001	0.1167	0.1237	0.0659	0.0735	0.0005	0.0005	0.0003	0.0003	0.0046
Inefficiency	0.0947	0.0368	0.0085	0.0248	0.0091	0.0202	0.0173	0.0168	0.0155	0.0056
Random error	0.0541	0.1096	0.1095	0.0324	0.0495	0.0306	0.0285	0.0280	0.0271	0.0177

5.2 In consideration of beta spillover effects

Recent studies address volatility spillover effects and commonalities in dynamic dependencies in multivariate volatility modeling (Herskovic et al., 2016; Bollerslev, et al., 2018). Similarly, we consider beta spillover effects and construct Vector HAR (VHAR) models to forecast betas. Specifically, the VHAR models can be regarded as a parsimonious version of the HAR-type model and are estimated by pool-fitting method (Bollerslev et al., 2018). To compare the performance of the competing models, we adopt R_{oos}^2 to evaluate the out-of-sample forecasting performance against the benchmark HAR-type model, as shown in Equation (24).

$$R_{OOS}^2 = 1 - \frac{\sum_{t=T_1+1}^{T} \omega_{i,t}\left(\beta_{i,t}^{real} - \widehat{\beta_{i,t}^{m}}\right)^2}{\sum_{t=T_1+1}^{T} \omega_{i,t}\left(\beta_{i,t}^{real} - \widehat{\beta_{i,t}^{HAR}}\right)^2}, \tag{24}$$

where $\beta_{i,t}^{real}$ refers to the actual realized beta of stock i at time t, $\widehat{\beta_{i,t}^{m}}$ and $\widehat{\beta_{i,t}^{HAR}}$ are the forecasted realized betas from model m and the HAR model, respectively, and $\omega_{i,t}$ is the weight of forecasting errors. Specifically, we consider two types of weighting methods to compute R_{OOS}^2:

equal-weight and value-weight. In the equal-weight scenario, $\omega_{i,t}$ is a constant $1/N$, while in the value-weight scenario, $\omega_{i,t}$ depends on the market capitalization of each stock.

The R^2_{OOS} results are shown in Table 9. We find that all VHAR models considering beta spillover effects have better forecast precision compared to the benchmark HAR model in both equal-weight and value-weight scenarios. The $\text{VHAR}_{\text{ARE.SB}}$ model, which includes indicators of both structural breaks and asymmetric risk effects, has the highest R^2_{OOS} among all VHAR models, with values of 10.84% and 10.81% for the two weighting scenarios, respectively. However, we find that the univariate $\text{LHAR}_{\text{ARE.SB}}$ model, implemented by the LASSO method, still outperforms all VHAR-type models with R^2_{OOS} values of 13.97% and 13.90% for the two weighting scenarios.

Table 9

Forecast performance considering beta spillover effects. This table presents the average R^2_{oos} of VHAR models that take beta spillover effects into consideration. Equal-weighted R^2_{oos} refers to R^2_{oos} based on equal weights of 1/*N*, while value-weighted R^2_{oos} refers to R^2_{oos} weighted by the market values of the stocks.

Models	Equal-weighted R^2_{oos}	Value-weighted R^2_{oos}
VHAR	4.28%	4.25%
VHAR_{ARE}	6.76%	6.74%
VHAR_{SB}	7.13%	7.10%
$\text{VHAR}_{\text{ARE.SB}}$	10.84%	10.81%
$\text{LHAR}_{\text{ARE.SB}}$	13.97%	13.90%

5.3 Alternative features and machine learning techniques

As shown in the literature, machine learning techniques contribute to the selection of predictors and improve predictive accuracy (Gu et al., 2020; Murray et al., 2022; Guijaro-Ordonez et al., 2022; Jiang et al., 2023; Bollerslev et al., 2023). Therefore, we introduce additional machine learning techniques, such as Ridge regression, Elastic-Net regression, and conventional principal component regression (PCR), to the HAR models with all the proposed predictors, including those for structural breaks and asymmetric risk effects. We compare the forecasting models using R^2_{oos} as shown in Eq.(24).

Moreover, to address the influence of structural breaks on the slopes of the HAR models, we introduce interaction predictors by interacting the break dummies with ARE predictors. These interaction terms consist of the structural break dummy variable and all other predictors, including the asymmetric risk effect predictors and the basic HAR predictors (daily, weekly and monthly realized betas), and we denote this model as $\text{HAR}_{\text{ARE.SB.interact}}$.

The results are shown in Table 10. Panel A and Panel B display the equal-weighted and value-weighted R_{oos}^2 results for forecasting models based on various machine learning methods. Table 10 suggests that forecasting models based on the LASSO method outperform those based on alternative machine learning methods and the benchmark PCR method. Moreover, considering interaction terms does not improve forecast accuracy but increases model complexity and computational burden.

Table 10

Forecast performance considering alternative features and machine learning techniques. This table presents the weighted average R_{oos}^2 of HAR-type models that consider asymmetric risk effects, structural break dummies, and interaction predictors constructed by interacting break dummies and other predictors. We employ alternative machine learning techniques such as Ridge regression, Elastic-Net regression, and PCR for predictor selection. Equal-weighted R_{oos}^2 refers to R_{oos}^2 based on equal weights of 1/*N*, while value-weighted R_{oos}^2 refers to R_{oos}^2 weighted by the market values of the stocks.

Panel A: Equal-weighted R_{oos}^2				
	LASSO	Ridge	ENet	PCR
HAR_{ARE}	10.27%	6.36%	7.94%	6.90%
$\text{HAR}_{\text{ARE.SB}}$	13.97%	11.91%	10.58%	9.06%
$\text{HAR}_{\text{ARE.SB.interact}}$	14.12%	12.08%	10.96%	10.57%
Panel B: Value-weighted R_{oos}^2				
	LASSO	Ridge	ENet	PCR
HAR_{ARE}	10.30%	6.52%	7.98%	6.91%
$\text{HAR}_{\text{ARE.SB}}$	13.90%	11.98%	10.79%	9.07%
$\text{HAR}_{\text{ARE.SB.interact}}$	14.18%	12.08%	11.03%	10.57%

5.4 The importance of predictors

This section analyzes how important the proposed new predictors are for forecasting realized betas and how their importance may vary over time. Figure 8 shows the time-series plots of frequencies at which daily, weekly and monthly asymmetric risk predictors (i.e., the four realized semi-betas and the two realized upside and downside betas) are selected by LASSO approaches. More specifically, given a rolling estimation window, we apply the LASSO method to $\text{HAR}_{\text{ARE.SB}}$ regression for each individual stock, the selection rate for each proposed new predictor is then calculated across all 327 individual stocks. It is observed that the selection rates for all asymmetric risk predictors vary substantially over time and it seems that no recognizable patterns emerge. The average selection rates over all rolling windows are reported in Table 11. We find that amongst all the daily and weekly predictors, realized semi-beta β^{M-} and realized downside beta β^{-} have the highest selection rates, respectively (see entries in the columns denoted by "daily" and "weekly"). In fact, the realized downside beta β^{-} predictors constructed over different horizons have very similar selection rates. This observation is also confirmed by time-series plots shown in Figure 8. The three β^{-} curves fluctuate in a relatively narrow range centered at around 0.87. The above findings indicate that the downside risk of financial markets has a significant impact on the beta forecasts. Another interesting finding is that the long-run positive risk component β^{+} is also an important predictor for realized betas. Finally, Figure 9 visualizes the time series of average number of structural breaks selected by the LASSO approach across all 327 individual stocks. It is observed that the number of structural breaks varies tightly between 4.7 and 5.2, which indicates that the structural break dummies also provide nontrivial predictive power for realized betas.

Table 11

Average selection rates for various components of realized beta. This table summarizes the average selection rates calculated over all estimation windows. See the notes to Figure 7 for more details.

	daily	weekly	monthly
β^{p}	0.8722	0.7921	0.6993
β^{M-}	0.8917	0.8414	0.7859
β^{M+}	0.8784	0.8192	0.7632
β^{N}	0.8730	0.8053	0.8220
β^{+}	0.8377	0.8438	0.9150
β^{-}	0.8879	0.8744	0.8617

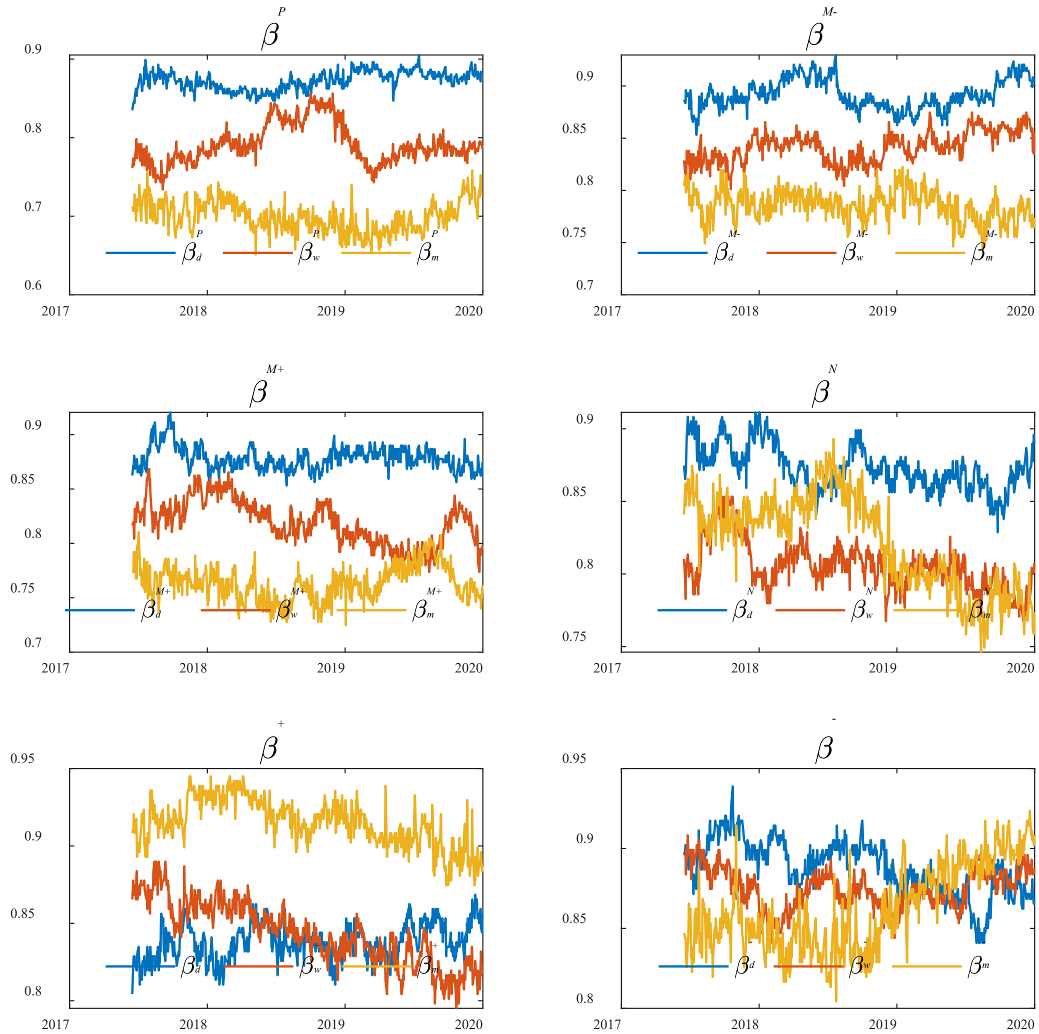

Figure 8. Selection rates for various components of realized beta. This figure shows the time-series plots of frequencies at which daily, weekly and monthly asymmetric risk predictors, including the four realized semi-betas and the two realized upside and downside betas, are selected by the LASSO approach over all rolling estimation windows.

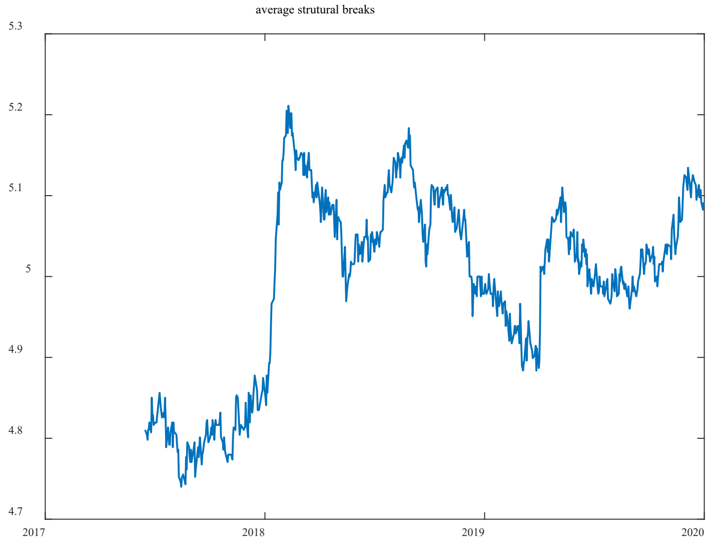

Figure 9. Average number of structural breaks. This figure shows the time-series plots of the average number of structural breaks across all individual stocks.

6. Robustness Check

6.1. Different loss functions

We consider two heteroscedasticity-adjusted statistics, the Heteroscedasticity-adjusted RMSE (HRMSE) and the Heteroscedasticity-adjusted MAE (HMAE) to better accommodate the heteroscedasticity in the forecast errors of realized betas (Andersen et al., 1999).

$$\text{HRMSE} = \sqrt{\frac{1}{\vartheta}\sum_{t=1}^{\vartheta}\left[\left(\beta_{i,t+h} - \hat{\beta}_{i,t+h}\right)/\beta_{i,t+h}\right]^2}, \tag{25}$$

$$\text{HMAE} = \frac{1}{\vartheta}\sum_{t=1}^{\vartheta}\left|\left(\beta_{i,t+h} - \hat{\beta}_{i,t+h}\right)/\beta_{i,t+h}\right|, \tag{26}$$

where ϑ is the total number of beta forecasts over the out-of-sample period, $\beta_{i,t+h}$ is the true value of realized beta on day t+h and $\hat{\beta}_{i,t+h}$ is the corresponding ex ante forecast.

The results are shown in Table A2 in the Appendix. The findings in general are still similar to those in Tables 3 and 4. Combining $\text{HAR}_{\text{ARE.SB}}$ with the LASSO method remains the best forecasting approach under both heteroscedasticity-adjusted loss functions.

6.2. Different rolling estimation windows

For robustness checks, we also replicate the forecast experiment using rolling windows of 400 days, 800 days, and 1000 days. The results are shown in Table A3. we find that the true long-memory models, including ARFIMA and FI, and the HAR variants with structural break dummies and asymmetric risk estimators have comparable forecasting performance. Our proposed HAR_{ARE}, HAR_{SB} and $HAR_{ARE.SB}$ models consistently generate more accurate beta forecasts than the classic HAR in terms of RMSE. Amongst all the competing models, $LHAR_{ARE.SB}$ still performs best.

6.3. Different data sampling frequencies

Moreover, we consider to use data sampled at alternative frequencies to construct various realized measures described in Section 2.3. With higher-frequency data, realized measures may be deteriorated due to larger microstructure noise. As a result, we follow Bollerslev et al. (2016a) and Becker et al. (2021) to consider the 30-minute sampling frequency in this section. We report the forecasting results for daily realized betas and monthly realized betas in Tables A5. The results are similar and consistent with baseline results as discussed above.

6.4. Different sub-periods

Moreover, we divide the sample period into three sub-periods: January 2007 to December 2011, January 2012 to December 2015, and January 2016 to December 2019, following Bollerslev et al. (2023). The first sub-period corresponds to the global financial crisis, the second to the market recovery period, and the last to the steady growth period. We specify the rolling windows for 1/5 of the full-sample length. The results, shown in Table A4, indicate that the proposed HAR_{ARE}, HAR_{SB} and $HAR_{ARE.SB}$ models exhibit better forecasting performance compared to other models. These three models particularly excel during the second sub-period when the market is in an upturn.

6.5. Different quantiles of the beta distribution

Additionally, we examine the forecasting performances of beta models across different quantiles of the beta distribution. Specifically, we sort the betas into quantiles and conduct a robust analysis for the lowest 1/5 betas, the middle betas, and the highest 1/5 betas. For each quintile of betas, the individual stocks are sorted into 10 portfolios. The results are shown in Table A6. We find that the $LHAR_{ARE.SB}$ model performs the best among the HAR-type models for all three scenarios. For the upper-tail betas, all HAR-type models outperform the conventional benchmark models, including RW, AR, ARMA, ARFIMA, and FI models. For the lower-tail betas, only the $LHAR_{ARE.SB}$ model outperforms the conventional ARMA, ARFIMA, and FI models in most cases. However, for the middle-beta sample, we find that HAR-type models perform worse than the conventional ARFIMA and FI models

6.6. Different sector-specific portfolios

Finally, we evaluate all the models based on forecasting results with sector-specific portfolios. We categorize all the 327 individual stocks and divide them into 12 different industry groups, including "Food Catering and Retail", "Energy and Electric", "Finance and Management", "Biomedicines", "Semiconductor and Electronic Components", "Industrial Products", "Information Technology Service", "Hotels, Entertainment and Media", "Automobile and Machinery", "Aerospace", "Bioscience" and "Others" (see Table A1 in Appendix A). Then we forecast realized betas of each portfolio. The results are shown in Table A7. With sector-specific portfolios, we still find that HAR_{SB}, $HAR_{ARE.SB}$ and $LHAR_{ARE.SB}$ models show the best performance. The classic HAR model is outperformed by other HAR variants, indicating, again, that the structural break dummies and asymmetric risk estimators provide nontrivial predictive power for realized betas. Therefore, the results are robust to the baseline results.

7. Conclusion

Beta is crucial to asset pricing, portfolio allocation and risk management, which can be consistently estimated by realized beta estimator using high-frequency financial data. Empirically, we find that realized betas have a number of prominent features, including long memory, structural breaks, and asymmetric risk effects. To characterize these features, we propose a set of new predictors for which structural breaks are detected by Iterated Cumulative Sum of Square (ICSS) algorithm and asymmetric risk effects are captured by decomposing the realized beta further into various components following Ang et al. (2006) and Bollerslev et al. (2021). The long-run dependence of realized betas can be reproduced by considering the HAR-type models that incorporate structural breaks and asymmetric risk effects. In addition, to achieve model parsimony and to keep only the predictors with significant power, we apply Least Absolute Shrinkage and Selection Operator (LASSO) method to the proposed HAR-type models for variable selection. We compare these new HAR-type models with a variety of competing models, including difference-stationary, short-memory and long-memory models through a comprehensive study of their forecasting performance.

The main empirical findings can be summarized as follow. (i) The difference-stationary and short-memory models are not suitable for forecasting realized betas; (ii) The true long-memory models show decent forecasting performance; (iii) The estimators of structural breaks and asymmetric risk effects provide nontrivial predictive power for realized betas; (iv) The LASSO approach further improves HAR-type models by dynamically selecting the most important predictors, which may vary substantially over time.

Declaration of interests

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

References

Andersen, T. G., Bollerslev, T., Diebold, F. X., & Labys, P. (2000). Great realizations. *Risk*, *13*, 105-108.

Andersen, T. G., Bollerslev, T., & Diebold, F. X. (2007). Roughing it up: Including jump components in the measurement, modeling, and forecasting of return volatility. *The review of economics and statistics*, 89(4), 701-720.

Andersen, T. G., Bollerslev, T., & Huang, X. (2011). A reduced form framework for modeling volatility of speculative prices based on realized variation measures. *Journal of econometrics*, 160(1), 176-189.

Andersen, T. G., Bollerslev, T., & Lange, S. (1999). Forecasting financial market volatility: Sample frequency vis-a-vis forecast horizon. *Journal of empirical finance*, 6(5), 457-477.

Andersen, T. G., Bollerslev, T., Diebold, F. X., & Wu, G. (2006). Realized beta: Persistence and predictability. *In Econometric analysis of financial and economic time series* (Vol. 20, pp. 1-39). Emerald Group Publishing Limited.

Andersen, T. G., Bollerslev, T., Diebold, F. X., & Wu, J. (2005). A framework for exploring the macroeconomic determinants of systematic risk. *American economic review*, 95(2), 398-404.

Andrews, D. W. (1991). Heteroskedasticity and autocorrelation consistent covariance matrix estimation. *Econometrica: journal of the econometric society*, 59(3), 817-858.

Ang, A., Chen, J., & Xing, Y. (2006). Downside risk. *The review of financial studies*, 19(4), 1191-1239.

Atilgan, Y., Bali, T. G., Demirtas, K. O., & Gunaydin, A. D. (2018). Downside beta and equity returns around the world. *The journal of portfolio management*, 44(7), 39-54.

Aye, G. C., Balcilar, M., Gupta, R., Kilimani, N., Nakumuryango, A., & Redford, S. (2014). Predicting

BRICS stock returns using ARFIMA models. *Applied financial economics*, 24(17), 1159-1166.

Baillie, R. T. (1996). Long memory processes and fractional integration in econometrics. *Journal of econometrics*, 73(1), 5-59.

Baillie, R. T., Calonaci, F., Cho, D., & Rho, S. (2019). Long memory, realized volatility and heterogeneous autoregressive models. *Journal of time series analysis*, 40(4), 609-628.

Barndorff-Nielsen, O. E., Hansen, P. R., Lunde, A., & Shephard, N. (2009). Realized kernels in practice: Trades and quotes. *Econometrics journal*, 12(3), 1–32.

Becker, J., Hollstein, F., Prokopczuk, M., & Sibbertsen, P. (2021). The memory of beta. *Journal of banking & finance*, 124, 106026.

Bhardwaj, G., & Swanson, N. R. (2006). An empirical investigation of the usefulness of ARFIMA models for predicting macroeconomic and financial time series. *Journal of econometrics*, 131(1-2), 539-578.

Bhattacharyya, M., Kumar M, D., & Kumar, R. (2009). Optimal sampling frequency for volatility forecast models for the Indian stock markets. Journal of Forecasting, 28(1), 38-54.

Bollerslev, T., & Mikkelsen, H. O. (1996). Modeling and pricing long memory in stock market volatility. *Journal of econometrics*, 73(1), 151-184.

Bollerslev, T., Li, S. Z., & Todorov, V. (2016). Roughing up beta: Continuous versus discontinuous betas and the cross section of expected stock returns. *Journal of financial economics*, 120(3), 464-490.

Bollerslev, T., Li, S. Z., Tang, Y., (2023). Forecasting and managing correlation risks. *Working paper*, Duke University and Rutgers University.

Bollerslev, T., Patton, A. J., & Quaedvlieg, R. (2016). Exploiting the errors: A simple approach for improved volatility forecasting. *Journal of Econometrics*, 192(1), 1-18.

Bollerslev, T., Hood, B., Huss, J., Pedersen, L. H., (2018). Risk everywhere: Modeling and managing volatility. *Review of Financial Studies*, 31, 2729–2773.

Bollerslev, T., Li, J., Patton, A. J. & Quaedvlieg. R. (2020). Realized semicovariances. *Econometrica*, 88(4), 1515-1551.

Bollerslev, T., Patton, A. J., & Quaedvlieg, R. (2021). Realized semibetas: Disentangling “good” and “bad” downside risks. *Journal of financial economics*, 144(1), 227-246.

Cenesizoglu, T., Ribeiro, F. D. O. F., & Reeves, J. J. (2017). Beta forecasting at long horizons. *International journal of forecasting*, 33(4), 936-957.

Choi, B. (1997). A recursive algorithm for solving the spatial Yule-Walker equations of causal spatial AR models. *Statistics & probability letters*, 33(3), 241-251.

Corsi, F. (2009). A simple approximate long-memory model of realized volatility. *Journal of financial econometrics*, 7(2), 174-196.

Corsi, F., Mittnik, S., Pigorsch, C., & Pigorsch, U. (2008). The volatility of realized volatility. *Econometric reviews*, 27(1-3), 46-78.

Daniel, K., Hirshleifer, D., & Sun, L. (2020). Short-and long-horizon behavioral factors. *The review of financial studies*, 33(4), 1673-1736.

Ding, Z., & Granger, C. W. (1996). Modeling volatility persistence of speculative returns: a new approach. *Journal of econometrics*, 73(1), 185-215.

Drobetz, W., Hollstein, F., Otto, T., & Prokopczuk, M. (2021). Estimating security betas via machine learning. *Available at SSRN.*

Epps, T. W. (1979). Comovements in stock prices in the very short run. *Journal of the American statistical association*, 74(366a), 291-298.

Fama, E. F. (1965). The behavior of stock-market prices. *The journal of business*, 38(1), 34-105.

Fama, E. F., & MacBeth, J. D. (1973). Risk, return, and equilibrium: Empirical tests. *Journal of political economy*, 81(3), 607-636.

Frazzini, A., & Pedersen, L. H. (2014). Betting against beta. *Journal of financial economics*, 111(1), 1-25.

Geweke, J., & Porter-Hudak, S. (1983). The estimation and application of long memory time series models. *Journal of time series analysis*, 4(4), 221-238.

Ghysels, E., & Jacquier, E. (2006). Market beta dynamics and portfolio efficiency. *Available at SSRN 711942*.

Granger, C. W., & Joyeux, R. (1980). An introduction to long-memory time series models and fractional differencing. *Journal of time series analysis*, 1(1), 15-29.

Gu, S., Kelly, B., & Xiu, D. (2020). Empirical asset pricing via machine learning. *The Review of Financial Studies*, *33*(5), 2223-2273.

Guijarro-Ordonez, J., Pelger, M., & Zanotti, G. (2021). Deep learning statistical arbitrage. *arXiv preprint arXiv:2106.04028*.

Hansen, P. R., Lunde, A., & Nason, J. M. (2011). The model confidence set. *Econometrica*, 79(2): 453-497.

Harvey, D., Leybourne, S., & Newbold, P. (1997). Testing the equality of prediction mean squared errors. *International Journal of forecasting*, 13(2), 281-291.

Herskovic, B., Kelly, B., Lustig, H., Van Nieuwerburgh, S., (2016). The common factor in idiosyncratic volatility: Quantitative asset pricing implications. *Journal of Financial Economics*, 119, 249–283.

Hollstein, F., & Prokopczuk, M. (2016). Estimating beta. *Journal of financial and quantitative analysis*, 51(4), 1437-1466.

Hollstein, F., Prokopczuk, M., & Simen, C. W. (2019). Estimating beta: Forecast adjustments and the impact of stock characteristics for a broad cross-section. *Journal of Financial Markets*, *44*, 91-118.

Hollstein, F., Prokopczuk, M., & Wese Simen, C. (2020). The conditional Capital Asset Pricing Model revisited: Evidence from high-frequency betas. *Management science*, 66(6), 2474-2494.

Hooper, V. J., Ng, K., & Reeves, J. J. (2008). Quarterly beta forecasting: An evaluation. *International journal of forecasting*, 24(3), 480-489.

Hosking, J. R. (1984). Modeling persistence in hydrological time series using fractional differencing. *Water resources research*, 20(12), 1898-1908.

Hosking, J. R. (1996). Asymptotic distributions of the sample mean, autocovariances, and autocorrelations of long-memory time series. *Journal of econometrics*, 73(1), 261-284.

Hosking, J., (1981). Fractional differencing. *Biometrika*, 68 (1), 165–176.

Inclan, C., & Tiao, G. C. (1994). Use of cumulative sums of squares for retrospective detection of changes of variance. *Journal of the American statistical association*, 89(427), 913-923.

Jiang, J., Kelly, B., & Xiu, D. (2023). (Re-) Imag (in) ing price trends. *The Journal of Finance*, *78*(6), 3193-3249.

Lettau, M., Maggiori, M., & Weber, M. (2014). Conditional risk premia in currency markets and other asset classes. *Journal of financial economics*, 114(2), 197-225.

Levi, Y., & Welch, I. (2017). Best practice for cost-of-capital estimates. *Journal of financial and quantitative analysis*, 52(2), 427-463.

Levi, Y., & Welch, I. (2020). Symmetric and asymmetric market betas and downside risk. *The review of financial studies*, 33(6), 2772-2795.

Liang, C., Xu, Y., Chen, Z., & Li, X. (2022). Forecasting China's stock market volatility with shrinkage method: Can Adaptive Lasso select stronger predictors from numerous predictors?. *International journal of finance & economics*, 1-11.

Linter, J. (1965). The valuation of risk assets and the selection of risky investments in stock portfolios and capital budgets. *Review of economics and statistics*, 47(1), 13-37.

McNeish, D. M. (2015). Using lasso for predictor selection and to assuage overfitting: A method long overlooked in behavioral sciences. *Multivariate behavioral research*, 50(5), 471-484.

Mincer, J. A., & Zarnowitz, V. (1969). The evaluation of economic forecasts. *In Economic forecasts and expectations: Analysis of forecasting behavior and performance* (pp. 3-46). NBER.

Murray, S., Xia, Y., & Xiao, H. (2024). Charting by machines. *Journal of Financial Economics*, *153*, 103791.

Papageorgiou, N., Reeves, J. J., & Xie, X. (2016). Betas and the myth of market neutrality. International *Journal of Forecasting*, 32(2), 548-558.

Patton A.J. & Verardo M. (2012). Does beta move with news? Firm-specific information flows and learning about profitability. *Review of financial studies*, 25(9), 2789-2839.

Post, T., & van Vliet, P. (2004). Do multiple factors help or hurt?. *Available at SSRN 582101*.

Reeves, J. J., & Wu, H. (2013). Constant versus Time-Varying Beta Models: Further Forecast Evaluation. *Journal of forecasting*, 32(3), 256-266.

Sagaert, Y. R., Aghezzaf, E. H., Kourentzes, N., & Desmet, B. (2018). Tactical sales forecasting using a very large set of macroeconomic indicators. *European Journal of Operational Research*, *264*(2), 558-569.

Scholes, M., & Williams, J. (1977). Estimating betas from nonsynchronous data. *Journal of financial economics*, 5(3), 309-327.

Sharpe, W. F. (1964). Capital asset prices: A theory of market equilibrium under conditions of risk. *The journal of finance*, 19(3), 425-442.

Siliverstovs, B. (2015). Short-term forecasting with mixed-frequency data: A MIDASSO approach. *KOF working papers*, no. 375.

Stock, J. H., & Watson, M. W. (2002). Macroeconomic forecasting using diffusion indexes. *Journal of business & economic statistics*, 20(2), 147-162.

Tofallis, C. (2008). Investment volatility: A critique of standard beta estimation and a simple way forward. *European Journal of Operational Research*, 187(3), 1358-1367.

Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Methodological)*, 58(1), 267-288.

Wang, K. Q. (2003). Asset pricing with conditioning information: A new test. *The journal of finance*, 58(1), 161-196.

Ziel, F., & Liu, B. (2016). Lasso estimation for GEFCom2014 probabilistic electric load forecasting. *International journal of forecasting*, 32(3), 1029-1037.

Zabarankin, M., Pavlikov, K., & Uryasev, S. (2014). Capital asset pricing model (CAPM) with drawdown measure. *European Journal of Operational Research*, *234*(2), 508-517.

Appendix A

Table A1

Industrial classification. This table summarizes the number of stocks for each industrial sector.

Industry Classification	Number of stocks
Food catering and retail	60
Energy and electric	43
Finance and management	38
Biomedicines	35
Semiconductor and electronic components	39
Industrial products	21
Information technology service	18
Hotels, entertainment and media	18
Automobile and machinery	17
Aerospace	12
Bioscience	11
Others	15
Total	**327**

Table A2

Out-of-sample forecasting results based on HRMSE and HMAE loss functions. This table shows the results of out-of-sample short-term (h = 1) prediction based on HRMSE (left of the forward slash) and HMAE (right of the forward slash) loss functions with rolling-window scheme. See the notes to Tables 3 for further details.

	RW	AR	ARMA	ARFIMA	FI	HAR	HAR_{ARE}	HAR_{SB}	$HAR_{ARE.SB}$	$LHAR_{ARE.SB}$
MCS_{80}	2/0	2/0	2/0	6/2	4/0	2/0	2/0	2/0	2/0	20/20
MCS_{90}	3/0	3/0	3/0	9/3	6/0	3/0	3/0	3/0	3/0	20/20
vs. RW	0/0	7/8	8/11	17/18	18/18	19/19	19/19	19/19	19/19	20/20
vs. AR	13/12	0/0	15/16	20/20	20/20	19/20	20/20	20/20	20/20	20/20
vs. ARMA	12/9	5/4	0/0	20/20	20/20	19/20	20/20	20/20	20/20	20/20
vs. ARFIMA	3/2	0/0	0/0	0/0	10/7	11/10	14/11	14/11	14/12	16/19
vs. FI	2/2	0/0	0/0	10/13	0/0	12/12	13/14	13/14	13/14	14/18
vs. HAR	1/1	1/0	1/0	9/10	8/8	0/0	19/19	20/20	20/20	20/20
vs. HAR_{ARE}	1/1	0/0	0/0	6/9	7/6	1/1	0/0	12/16	19/20	19/20
vs. HAR_{SB}	1/1	0/0	0/0	6/9	7/6	0/0	8/4	0/0	19/19	20/20
vs. $HAR_{ARE.SB}$	1/1	0/0	0/0	6/8	7/6	0/0	1/0	1/1	0/0	19/20
vs. $LHAR_{ARE.SB}$	0/0	0/0	0/0	4/1	6/2	0/0	1/0	0/0	1/0	0/0

Table A3

Forecasting results based on shorter rolling estimation windows. This table shows the results of out-of-sample short-term (h = 1) prediction using 400-day (left of the forward slash), 800-day (middle of the forward slash) and 1000-day (right of the forward slash) rolling estimation windows. Predictive accuracy is measured by the RMSE loss function. Refer to the notes in Table 3 for more details.

	RW	AR	ARMA	ARFIMA	FI	HAR	HAR_{ARE}	HAR_{SB}	$HAR_{ARE.SB}$	$LHAR_{ARE.SB}$
MCS_{80}	0/0/0	0/0/0	0/0/0	7/7/7	1/1/1	0/0/0	0/0/0	0/0/0	0/0/0	17/17/17
MCS_{90}	0/0/0	0/0/0	0/0/0	8/8/8	2/1/1	0/0/0	0/0/0	0/0/0	0/0/0	17/18/18
vs. RW	0/0/0	15/14/15	16/15/15	17/17/17	17/17/17	20/20/20	20/20/20	20/20/20	20/20/20	20/20/20
vs. AR	5/6/5	0/0/0	19/19/19	20/20/20	20/20/20	18/19/19	18/19/19	19/19/19	19/19/19	19/19/19
vs. ARMA	4/5/5	1/1/1	0/0/0	20/20/20	19/20/20	16/16/16	16/17/17	16/17/17	16/17/17	18/18/19
vs. ARFIMA	3/3/3	0/0/0	0/0/0	0/0/0	7/6/6	6/7/7	7/9/8	7/9/8	8/9/9	12/12/12
vs. FI	3/3/3	0/0/0	1/0/0	13/14/14	0/0/0	9/9/9	10/10/11	10/10/11	10/12/11	15/14/14
vs. HAR	0/0/0	2/1/1	4/4/4	14/13/13	11/11/11	0/0/0	20/20/20	20/20/20	20/20/20	20/20/20
vs. HAR_{ARE}	0/0/0	2/1/1	4/3/3	13/11/12	10/10/9	0/0/0	0/0/0	17/19/19	20/20/20	20/20/20
vs. HAR_{SB}	0/0/0	1/1/1	4/3/3	13/11/12	10/10/9	0/0/0	3/1/1	0/0/0	20/20/20	20/20/20
vs. $HAR_{ARE.SB}$	0/0/0	1/1/1	4/3/3	12/11/11	10/8/9	0/0/0	0/0/0	0/0/0	0/0/0	20/20/20
vs. $LHAR_{ARE.SB}$	0/0/0	1/1/1	2/2/1	8/8/8	5/6/6	0/0/0	0/0/0	0/0/0	0/0/0	0/0/0

Table A4

Forecasting results based on prior sub-periods rolling estimation windows. This table shows the results of out-of-sample short-term (h = 1) prediction using the sub-periods from January 2007 to December 2011 (left of the forward slash), January 2012 to December 2015 (middle of the forward slash) and January 2016 to December 2019 (right of the forward slash). Predictive accuracy is measured by the RMSE loss function. See the notes to Table 3 for more details.

	RW	AR	ARMA	ARFIMA	FI	HAR	HAR_{ARE}	HAR_{SB}	$HAR_{ARE.SB}$	$LHAR_{ARE.SB}$
MCS_{80}	0/0/0	0/0/0	0/0/0	6/7/9	4/0/0	0/0/0	0/0/0	0/0/0	0/0/0	17/19/15
MCS_{90}	0/0/0	0/0/0	0/0/0	6/9/9	4/1/0	0/0/0	0/0/0	0/0/0	0/0/0	17/20/16
vs. RW	0/0/0	14/13/13	15/15/14	16/17/17	16/17/17	20/20/20	20/20/20	20/20/20	20/20/20	20/20/20
vs. AR	6/7/7	0/0/0	19/19/20	20/20/20	20/20/20	19/19/18	19/19/18	19/19/18	19/19/18	19/20/20
vs. ARMA	5/5/6	1/1/0	0/0/0	20/20/20	19/20/20	15/16/16	16/17/18	16/17/18	16/18/18	17/19/18
vs. ARFIMA	4/3/3	0/0/0	0/0/0	0/0/0	9/5/4	7/6/6	8/6/7	8/8/6	10/8/9	12/12/11
vs. FI	4/3/3	0/0/0	1/0/0	11/15/16	0/0/0	8/9/8	10/10/9	10/10/9	11/10/9	14/15/16
vs. HAR	0/0/0	1/1/2	5/4/4	13/14/14	12/11/12	0/0/0	20/20/20	20/20/20	20/20/20	20/20/20
vs. HAR_{ARE}	0/0/0	1/1/2	4/3/2	12/14/13	10/10/11	0/0/0	0/0/0	15/18/19	20/20/20	20/20/20
vs. HAR_{SB}	0/0/0	1/1/2	4/3/2	12/12/14	10/10/11	0/0/0	5/2/1	0/0/0	20/20/20	20/20/20
vs. $HAR_{ARE.SB}$	0/0/0	1/1/2	4/2/2	10/12/11	9/10/11	0/0/0	0/0/0	0/0/0	0/0/0	19/20/20
vs. $LHAR_{ARE.SB}$	0/0/0	1/0/0	3/1/2	8/8/9	6/5/4	0/0/0	0/0/0	0/0/0	1/0/0	0/0/0

Table A5

Short-term forecasting results based on data sampled at the 30-minute frequency. This table shows the results of out-of-sample short-term (h = 1) (left of the forward slash) and long-term (h = 22) (right of the forward slash) prediction using data sampled at the 30-minute frequency. Predictive accuracy is measured by the RMSE loss function. See the notes to Table 3 for more details.

	RW	AR	ARMA	ARFIMA	FI	HAR	HAR_{ARE}	HAR_{SB}	$HAR_{ARE.SB}$	$LHAR_{ARE.SB}$
MCS_{80}	0/0	0/0	0/0	7/12	1/3	0/0	0/0	0/0	0/0	17/11
MCS_{90}	0/0	0/0	0/0	8/12	1/3	0/0	0/0	0/0	0/0	18/11
vs. RW	0/0	15/14	15/15	17/19	17/19	20/20	20/20	20/20	20/20	20/20
vs. AR	5/6	0/0	19/19	20/20	20/20	19/18	19/18	19/19	19/19	19/19
vs. ARMA	5/5	1/1	0/0	20/20	20/20	16/16	17/16	17/16	17/16	18/18
vs. ARFIMA	3/1	0/0	0/0	0/0	6/3	7/4	8/4	8/4	9/4	12/6
vs. FI	3/1	0/0	0/0	14/17	0/0	9/4	11/5	11/5	12/5	14/8
vs. HAR	0/0	1/2	4/4	13/16	11/16	0/0	20/20	20/20	20/20	20/20
vs. HAR_{ARE}	0/0	1/2	3/4	12/16	9/15	0/0	0/0	19/16	20/20	20/20
vs. HAR_{SB}	0/0	1/1	3/4	12/16	9/15	0/0	1/4	0/0	20/20	20/20
vs. $HAR_{ARE.SB}$	0/0	1/1	3/4	11/16	8/15	0/0	0/0	0/0	0/0	20/20
vs. $LHAR_{ARE.SB}$	0/0	1/1	2/2	8/14	6/12	0/0	0/0	0/0	0/0	0/0

Table A6

Forecasting results based on beta distribution. This table shows the results of out-of-sample short-term (h = 1) prediction of lower-betas (left of the forward slash), middle-betas (middle of the forward slash) and higher-betas (right of the forward slash) based on the three quantiles of the beta distribution equally. Predictive accuracy is measured by the RMSE loss function. See the notes to Table 3 for more details.

	RW	AR	ARMA	ARFIMA	FI	HAR	HAR_{ARE}	HAR_{SB}	$HAR_{ARE.SB}$	$LHAR_{ARE.SB}$
MCS_{80}	0/0/0	0/0/0	0/0/0	0/10/0	0/0/0	0/0/0	0/0/0	0/0/0	0/0/0	10/0/10
MCS_{90}	0/0/0	0/0/0	0/0/0	0/10/0	0/0/0	0/0/0	0/0/0	0/0/0	0/0/0	10/1/10
vs. RW	0/0/0	6/10/4	7/10/5	10/10/8	10/10/8	10/10/10	10/10/10	10/10/10	10/10/10	10/10/10
vs. AR	4/0/6	0/0/0	10/10/10	10/10/10	10/10/10	10/8/10	10/8/10	10/8/10	10/8/10	10/9/10
vs. ARMA	3/0/5	0/0/0	0/0/0	10/10/10	10/10/10	10/0/10	10/1/10	10/1/10	10/1/10	10/8/10
vs. ARFIMA	0/0/2	0/0/0	0/0/0	0/0/0	10/0/2	3/0/9	4/0/9	5/0/9	5/0/9	6/0/9
vs. FI	0/0/2	0/0/0	0/0/0	0/10/8	0/0/0	1/0/9	2/0/9	3/0/9	3/0/9	5/2/10
vs. HAR	0/0/0	0/2/0	0/10/0	7/10/1	9/10/1	0/0/0	10/10/10	10/10/10	10/10/10	10/10/10
vs. HAR_{ARE}	0/0/0	0/2/0	0/9/0	6/10/1	8/10/1	0/0/0	0/0/0	8/9/8	10/10/10	10/10/10
vs. HAR_{SB}	0/0/0	0/2/0	0/9/0	5/10/1	7/10/1	0/0/0	2/1/2	0/0/0	9/10/10	10/10/10
vs. $HAR_{ARE.SB}$	0/0/0	0/2/0	0/9/0	5/10/1	7/10/1	0/0/0	0/0/0	1/0/0	0/0/0	10/10/10
vs. $LHAR_{ARE.SB}$	0/0/0	0/1/0	0/2/0	4/10/1	5/8/0	0/0/0	0/0/0	0/0/0	0/0/0	0/0/0

Table A7

Forecasting results based on sector-specific portfolios. This table shows the results of out-of-sample short-term ($h = 1$) forecasts of sector-specific portfolios' betas. Predictive accuracy is measured by the RMSE loss function. See the notes to Table 3 for more details. The construction of these portfolios is described in Section 6.3. See also Table A1 in Appendix A for details.

	RW	AR	ARMA	ARFIMA	FI	HAR	HAR_{ARE}	HAR_{SB}	$HAR_{ARE.SB}$	$LHAR_{ARE.SB}$
MCS80	0	0	0	3	0	0	0	0	0	11
MCS90	0	0	0	4	0	0	0	0	0	11
vs. RW	0	7	8	9	9	12	12	12	12	12
vs. AR	5	0	12	12	12	11	11	11	11	11
vs. ARMA	4	0	0	12	12	10	10	10	10	11
vs. ARFIMA	3	0	0	0	5	6	6	6	6	8
vs. FI	3	0	0	7	0	6	9	9	9	9
vs. HAR	0	1	2	6	6	0	12	12	12	12
vs. HAR_{ARE}	0	1	2	6	3	0	0	12	12	12
vs. HAR_{SB}	0	1	2	6	3	0	0	0	12	12
vs. $HAR_{ARE.SB}$	0	1	2	6	3	0	0	0	0	12
vs. $LHAR_{ARE.SB}$	0	1	1	4	3	0	0	0	0	0