The Performance of Characteristic-Sorted Portfolios: Evaluating the Past and Predicting the Future

AYDOĞAN ALTI, TRAVIS L. JOHNSON, and SHERIDAN TITMAN*

November 2021

Abstract

The expected returns of most portfolios are likely to fluctuate over time. We present a statistical model that allows for such fluctuations and apply the model to analyze the returns of characteristic-sorted portfolios, such as value minus growth. We find that accounting for plausible magnitudes of persistent variation in returns doubles the standard errors of these portfolios' expected return estimates. We also analyze characteristic-sorted portfolios from the perspective of Bayesian investors and show that investors' posterior beliefs about expected returns are highly dependent on their priors about persistence.

*All three authors are with the University of Texas at Austin. We thank seminar participants at Peking University and the University of Texas at Austin for helpful comments.

A large and growing literature links firm characteristics, such as valuation ratios, to expected rates of stock returns. While the evidence documented in this literature convincingly rejects the CAPM, going beyond this rejection and interpreting alternatives has proven to be challenging. For instance, some recent studies document time-variation in characteristic-sorted portfolio returns, and others note longer-term fluctuations, such as the weakening of the value premium or the increased performance of profitability-based strategies in recent decades.[1] A natural question arising from these observations is whether the historical links between characteristics and returns represent permanent economic forces that will continue to shape returns in the future, or transitory forces that will eventually dissipate.

To address this and related questions, we examine the returns of characteristic-sorted portfolios through the lens of a statistical model that allows expected returns to vary over time. The central feature of our model, persistent variation in expected returns, is implicit in existing rational and behavioral theories of characteristic-based return predictability. The general idea is that the economic forces that generate relationships between characteristics and returns are likely to change over time.[2] Our main contribution is to analyze how such persistent variation influences our inferences about the past as well as the predictions about the future returns of characteristic-sorted portfolios.

Our analysis delivers three main findings. First, allowing for plausible magnitudes of persistent variation in conditional expected returns has a large effect on the precision of unconditional expected return estimates, often doubling their standard errors. Second, the

[1] See Gupta and Kelly (2019) and Arnott et al. (2021a) for evidence of positive autocorrelation in characteristic-sorted portfolio returns, Fama and French (2021) for the weakening value effect, and Novy-Marx (2013) for the strengthening profitability effect.

[2] Gârleanu, Kogan, and Panageas (2012) and Kogan, Papanikolaou, and Stoffman (2020) provide examples of rational theories where investors hold lower returning growth stocks to hedge technology shocks. Shiller (2000) provides a behavioral explanation based on investors being overly optimistic about the commercialization potential of new technologies, causing growth stocks to underperform. One might expect the economic fundamentals that drive both rational and behavioral explanations to fluctuate over time, or possibly even reverse, causing growth stocks to outperform value stocks. See Altı and Titman (2019) for a dynamic behavioral model that captures both possibilities.

degree of persistent variation in expected returns is very imprecisely estimated. These two findings imply that the historical return data we analyze is consistent with very different return generating processes. For instance, the historical value portfolio returns are consistent with a persistent process with zero unconditional expected returns as well as an i.i.d. process with large unconditional expected returns. Given this wide range, it is natural to ask how Bayesian analysts with different prior beliefs interpret the data and forecast future return performance of characteristic-sorted portfolios. This leads to our third finding, which is that the posterior beliefs of Bayesian analysts are highly sensitive to priors about the degree of persistence, even after observing 56 years of data.

Our empirical analysis focuses on the returns of the four characteristic-sorted portfolios described in Fama and French (2015) – value, investment, profitability, and size. As a first step, we examine the autocorrelation patterns in these portfolios' returns. All four portfolios exhibit positive return autocorrelation over yearly horizons, and while only the size portfolio's autocorrelation estimate is significant on its own, a joint test of the four portfolios strongly rejects the zero-autocorrelation null hypothesis. These autocorrelation estimates, however, are sufficiently imprecise that they are consistent with a wide variety of parameterizations of our model, from persistence lasting months to decades.

Next, we apply our model to analyze the returns of the four characteristic-sorted portfolios. As we show, different assumptions about the extent of persistent variation in conditional expected returns, all consistent with the historical return data, generate substantially different inferences about unconditional expected returns. If returns are assumed to be i.i.d., meaning they have no persistence, standard errors are low and the hypothesis of zero unconditional expected returns is strongly rejected. However, if shocks to conditional expected returns are assumed to be somewhat persistent, both the ordinary least squares (OLS) and generalized least squares (GLS) standard errors are substantially higher and the hypothesis is not rejected. For example, the standard errors for all four strategies' unconditional expected return estimates more than double, relative to the case where returns are i.i.d., when the half-life of shocks to conditional expected returns are five years. Intuitively, when expected returns are persistent, the return time series exhibits less independent variation, generating less precise estimates of the unconditional expected returns.[3]

We also use maximum likelihood to estimate a version of our model where the degree of persistence is an estimated parameter, rather than assumed. Consistent with the highly noisy autocorrelation estimates in our reduced-form regressions, we find that the persistence parameter is imprecisely estimated. Given these estimates, a variety of a plausible alternative explanations for the historical performance of the characteristic-sorted portfolios cannot be ruled out. Indeed, for several of the portfolios, likelihood ratio tests fail to reject the possibility that the unconditional expected return is zero, but conditional expected returns are both variable and highly persistent.

While the results discussed so far are expressed from a frequentist statistical perspective, our focus on how assumptions about time-variation affect inferences about unconditional expected returns has a natural Bayesian interpretation. Specifically, our OLS and GLS estimates are similar to Bayesian analyses with dogmatic priors about the degree of persistent variation in expected returns, while the maximum likelihood estimations are similar to Bayesian analyses with agnostic priors. This analogy is limited, however, to priors about persistence, as the frequentist analysis always assumes fully agnostic priors about unconditional expected returns.

To assess how priors about unconditional expected returns and expected return persistence interact, we embed the model into a Bayesian framework. Specifically, we assume that different analysts start with different, but economically plausible priors, and that they update their beliefs based on the same return data that we observe in our sample period. The

[3]Formally and more generally, when error terms are positively serially correlated, true standard errors are typically larger than their OLS estimates. See, for instance, Greene (2000).

value of this exercise is to quantify the extent to which individuals with different perspectives on the determinants of expected returns may differ in their posterior beliefs after observing the historical data. For example, analysts guided by predictions of the CAPM may have the prior belief that the unconditional expected returns of characteristic-sorted portfolios cannot deviate substantially from zero, while others may have less confidence in the CAPM.

We find that prior beliefs about persistence substantially affect how analysts, after observing the return history, update their beliefs about unconditional expected returns. If analysts have strong priors that expected returns fluctuate very little, then their posterior beliefs about unconditional expected returns tend to be relatively precise and not highly sensitive to priors about unconditional expected returns. However, if analysts' priors put more weight on the possibility of persistent fluctuations, then their posterior beliefs about unconditional expected returns become more diffuse and more sensitive to their priors. In this way, persistence makes analysts learn less from data and rely more on their priors.

We also calculate Bayesian analysts' estimates of the conditional expected returns and Sharpe ratios of the four characteristic-sorted portfolios at each point in time throughout the sample period.[4] These estimates illustrate the extent to which the analysts believe that the characteristic-sorted portfolio returns can be timed. The estimated annualized conditional Sharpe ratios exhibit substantial variation over time, ranging from lows around 0.1 to highs around 0.8 for the value, investment, and profitability portfolios. Interestingly, the size portfolio's conditional Sharpe ratio exhibits even greater and more persistent swings – peaking near 1.0 in 1968, 1979, and 2002, while also falling below -0.5 in 1988, 1999, and 2019 – despite the average conditional Sharpe ratio being close to zero. These large and persistent variations that the size portfolio returns exhibit have been largely unexplored in the literature.

[4]Specifically, we compute the means of an analyst's posteriors about the conditional expected returns and Sharpe ratios given their priors and data observations from the full sample period.

Finally, we use our Bayesian approach to generate estimates of conditional expected returns and Sharpe ratios in 2020, the year after our sample period ends. These estimates speak directly to the ongoing debate about the recent performance of value strategies and the effectiveness of factor timing strategies.[5] We find that Bayesian analysts who account for time-varying expected returns will have very different views of near-term and long-term expected returns. The reason is that while data from the early and later parts of the sample are equally important for estimating unconditional expected returns, near-term conditional expected return estimates put more weight on the more recent observations when returns are persistent. For example, because the value portfolio performed particularly poorly towards the end of the sample, the value premium in 2020, measured as the mean of the posterior, is generally around a quarter of the unconditional value premium, and is close to zero in many specifications. The profitability portfolio exhibits the opposite pattern: because the returns were stronger in recent decades, the posteriors for conditional expected returns in 2020 are higher than posteriors for the unconditional expected return.

While we are among the first to analyze time variation in the expected returns of characteristic-sorted portfolios, there is a well-established literature that explores a number of these issues within the context of the aggregate equity market portfolio. For example, Ferson, Sarkissian, and Simin (2003) consider the predictability of aggregate market returns using persistent predictor variables such as price/dividend ratios. They present simulations that show that OLS regressions can overstate the significance of such relationships in finite samples when expected returns are persistent, even when the Newey and West (1987) standard errors adjustment is used. Although our application is different, the problem we address is similar and we also show that the Newey and West (1987) correction is not effective in dealing with the problem. In addition, we propose some frequentist remedies, such

[5] See Arnott et al. (2021b), Asness et al. (2021), Choi, So, and Wang (2021), Eisfeldt, Kim, and Papanikolaou (2021), Fama and French (2021), and Goncalves and Leonard (2021).

as GLS regressions when the level of persistence in expected returns is known, or maximum likelihood when the level of persistence is estimated.

Pástor and Stambaugh (2009, 2012) also conduct Bayesian analyses of time-varying expected returns and find that the priors about the return generating process substantially affect the posteriors about expected returns, a similar conclusion to the one we reach. Our study differs from these papers in both application and focus, examining returns of characteristic-sorted portfolios instead of timing the aggregate equity market portfolio based on imperfect predictors.[6] We are aware of only one study, Pástor (2000), that uses Bayesian methods to study characteristic-sorted portfolio returns. However, in contrast to our analysis, persistence plays no role in Pástor (2000) as expected returns are assumed to be constant.

Finally, as we mentioned at the outset, our analysis builds on a number of recent papers that combine evidence from many characteristic-based portfolios to show that there is persistent time-variation in conditional expected returns, e.g., Lewellen (2002), McLean and Pontiff (2016), Avramov et al. (2017), Gupta and Kelly (2019), Arnott et al. (2021a), and Ehsani and Linnainmaa (2021). We contribute to this literature by studying the implications of this persistence. In addition, we find that while there is strong evidence of persistence when evaluating the returns of many portfolios, the magnitude of persistence for an individual portfolio is estimated so imprecisely that in most cases we fail to reject everything from zero autocorrelation and economically-important autocorrelations persisting for decades.

The remainder of the paper is organized as follows. Section 1 describes our statistical model of the return-generating process. Section 2 documents the historical performance and the evidence of return autocorrelation for the characteristic-sorted portfolios. Section 3 presents the OLS, GLS, and maximum likelihood estimations of the model. Section 5 presents the Bayesian analysis. Section 6 concludes.

[6] Other papers in the literature that employ Bayesian methods to study aggregate equity market returns include Kandel and Stambaugh (1996), Barberis (2000), Wachter and Warusawitharana (2009), and Johannes, Korteweg, and Polson (2014).

1. Statistical Model

In this section, we describe and analyze a simple statistical model of the return-generating process that we later apply to characteristic-sorted portfolio returns. The important feature of the model is the assumption that the conditional expected returns of the portfolio exhibits persistent variations around an unconditional mean.

Specifically, we assume that the time-series of the zero-cost portfolio returns r_t satisfies:

$$r_{t+1} = \mu_t + \epsilon_{t+1}, \tag{1}$$

$$\mu_{t+1} = \mu + \lambda(\mu_t - \mu + \delta_{t+1}), \tag{2}$$

where μ_t and μ are the conditional and unconditional expected returns, respectively, and $\lambda \geq 0$ is a parameter that determines the persistence of shocks to μ_t.[7] The shocks ϵ_t and δ_t are i.i.d. and follow a joint normal distribution with variances σ_ϵ and σ_δ, respectively, and correlation $\rho \in (-1, 1)$. We expect ρ to be negative, since shocks to expected rates of return, ceteris paribus, reduce an investment's value.

The econometrician does not observe μ_t, but can estimate it – along with other model parameters – from the observed return realizations $\mathbf{R} = [r_1, r_2, \ldots, r_T]'$. Conditional on parameters $\Omega = [\mu, \lambda, \sigma_\epsilon, \sigma_\delta, \rho,]$, $\mathbf{R}$ has the following mean and covariance matrix:

$$\mathbb{E}(\mathbf{R}|\Omega) = \mu, \tag{3}$$

$$\text{Cov}\,(\mathbf{R}|\Omega) = \mathbf{\Sigma}(\Omega), \qquad \mathbf{\Sigma}(\Omega)_{i,j} = \begin{cases} \frac{\lambda^2\sigma_\delta^2}{1-\lambda^2} + \sigma_\epsilon^2 & \text{if } i = j \\ \lambda^{|i-j|}\left(\frac{\lambda^2\sigma_\delta^2}{1-\lambda^2} + \rho\sigma_\delta\sigma_\epsilon\right) & \text{if } i \neq j \end{cases}. \tag{4}$$

[7] We multiply the shock δ_{t+1} by λ so that returns are i.i.d. when $\lambda = 0$.

Equation (4) shows that shocks to both the expected and unexpected returns contribute to the volatility of returns (the terms σ_δ^2 and σ_ϵ^2, respectively). Note also that the covariance between r_i and r_j decays at constant rate λ as $|i-j|$ grows.

The model is over-parameterized in the sense that multiple values of Ω lead to the same predicted moments $\mathbb{E}(\mathbf{R}|\Omega)$ and $\mathbf{\Sigma}(\Omega)$. To see this, define:

$$\sigma_r^2(\Omega) = \mathrm{Var}(r_t) = \frac{\lambda^2\sigma_\delta^2}{1-\lambda^2} + \sigma_\epsilon^2, \tag{5}$$

$$\gamma(\Omega) = \mathrm{Corr}(r_{t+1}, r_t) = \lambda\frac{\lambda^2\sigma_\delta^2 + (1-\lambda^2)\rho\sigma_\delta\sigma_\epsilon}{\lambda^2\sigma_\delta^2 + (1-\lambda^2)\sigma_\epsilon^2}. \tag{6}$$

Using this alternative notation, the covariance matrix becomes:

$$\mathbf{\Sigma}(\Omega)_{i,j} = \begin{cases} \sigma_r^2 & \text{if } i = j, \\ \lambda^{|i-j|-1}\gamma\sigma_r^2 & \text{if } i \neq j. \end{cases} \tag{7}$$

Note that any two parameterizations Ω and $\tilde{\Omega}$ satisfying

$$[\mu, \lambda, \sigma_r, \gamma] = \left[\tilde{\mu}, \tilde{\lambda}, \tilde{\sigma_r}, \tilde{\gamma}\right] \tag{8}$$

result in identical mean and covariance matrix for $\mathbf{R}$ (μ and $\mathbf{\Sigma}$). Because we have five parameters to satisfy four equations, many distinct Ω and $\tilde{\Omega}$ generate any given μ and $\mathbf{\Sigma}$.

To better understand identification in our model, note that the sample mean of returns identifies the unconditional expected return μ (Equation (3)), while the rate at which the covariance between r_i and r_j decays as $|i-j|$ grows identifies the persistence parameter λ (Equation (4)). This leaves three other parameters to be identified, σ_ϵ, σ_δ, and ρ, but only two other moments that can be estimated, the variance of returns σ_r^2 in Equation (5) and the one-lag autocorrelation γ in Equation (6).[8] Intuitively, the identification problem arises because one cannot distinguish between different channels that generate return variance and autocorrelation. An increase in the volatility of expected return shocks σ_δ increases both return variance and autocorrelation. But the same magnitudes of increases in these two moments can also be generated from increases in the volatility of unexpected return shocks σ_ϵ and the correlation parameter ρ.

We address this identification problem in our frequentist analysis by estimating the four moments $\theta = [\mu, \lambda, \sigma_r, \gamma]$, which we can identify, rather than the full set of underlying parameters Ω. In doing so, we apply the constraint that there must be a parameterization Ω which is consistent with θ and satisfies $\sigma_\epsilon > 0$, $\sigma_\delta > 0$, $\lambda \geq 0$, and $\rho \in (-1, 1)$. Our frequentist analysis does estimate Ω or assume anything about which Ω generates the θ we estimate. The identification problem is not an issue for our Bayesian analysis because we compute a posterior distribution for parameter values, which is unique given a set of priors and observed data, rather than a single point estimate, which is not unique.

To facilitate our interpretation of economic magnitudes, we express the persistence parameter λ as an annualized half-life H hereafter:

$$H = \frac{\log(0.5)}{\log(\lambda)} \frac{1}{N}, \tag{9}$$

where N is the number of periods per year (e.g. 4 for quarterly data).

2. Characteristic-Sorted Portfolios

We apply our statistical model to study four portfolios that are formed by sorting stocks based on firm characteristics: value, investment, profitability, and size. We focus on these portfolios because they are the basis of the Fama and French (2015) five-factor model and

[8]Note from Equation (7) that return covariances at longer lags do not provide any additional information about the model parameters, because all these covariance terms are scaled by $\gamma\sigma_r^2$.

illustrate a variety of channels through which time-varying expected returns affect our understanding of portfolio returns.

2.1. Data and Characteristic Definitions

We use data on historical returns of characteristic-sorted portfolios from Ken French's website.[9] Each portfolio combines a long position in a value-weighted portfolio of firms in one extreme quintile of the characteristic with a short position in the other extreme.

The characteristics are defined following Fama and French (2015). *Value* is the ratio of the book value of equity ($B_{i,y}$) to the market value of equity ($M_{i,y}$) as of the end of the prior fiscal year y. *Investment* is the growth rate in the book value of assets ($\text{Assets}_{i,y}/\text{Assets}_{i,y-1}$). *Profitability* is revenues minus cost of goods sold, interest expense, and selling, general, and administrative expenses in year y divided by book equity in year $y-1$. *Size* is the market value of equity $M_{i,y}$.

In contrast to most of the literature, which examines the monthly returns of characteristic sorted portfolios, we analyze quarterly returns. As we show in the Appendix, the monthly returns of three of the four portfolios we study exhibit strong positive first-order autocorrelations.[10] While these short-term autocorrelations are consistent with time-varying expected returns, they could also be driven by lead-lag effects and other short-term microstructure effects. For this reason, and more generally to focus our analysis on longer-term autocorrelations driven by persistent variations in expected returns, we estimate our model on quarterly returns for each portfolio, which have weaker first-order autocorrelations.

To analyze expected returns associated with each characteristic-sorted portfolio rather than the market risk premium, we use market-neutral versions of each portfolio's returns

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

[10] This is consistent with the evidence in Gupta and Kelly (2019) that 47 of 65 characteristic-based portfolios have significantly positive first-order autocorrelation.

throughout, calculated as:

$$r_{i,t}^{\beta=0} = r_{i,t} - \hat{\beta}_i(r_{m,t} - r_{f,t}), \tag{10}$$

were $r_{i,t}$ is the quarterly return of long-short portfolio i, $r_{m,t} - r_{f,t}$ is the quarterly excess market return, and $\hat{\beta}_i$ is the full-sample market beta. The average return and the Sharpe ratio of this market-hedged portfolio are equivalent to the alpha and information ratio, respectively, of the unhedged long-short portfolio.

2.2. Historical Performance of Characteristic-Sorted Portfolios

Table 1 summarizes the historical performance of the characteristic-sorted portfolios. Panel A shows that value, investment, and profitability portfolios all have annualized excess returns around 4% and Sharpe Ratios between 0.3 and 0.5. We calculate standard errors for these statistics under the assumption that returns are i.i.d. by taking the standard deviation across simulated samples formed by re-sampling historical data with replacement. Under this assumption, we can strongly reject the hypothesis that average returns equal zero, with t-statistics of 2.4, 3.3, and 3.5 for value, investment, and profitability, respectively.

We also examine the historical performance of the portfolios in the first and second halves of our sample, 1963–1990 and 1991–2019, respectively. As discussed in McLean and Pontiff (2016), Linnainmaa and Roberts (2018), and Fama and French (2021), the value strategy's average returns are smaller in the second half of the sample and not statistically different from zero. However, as emphasized by Fama and French (2021), the difference between the two halves of the sample is not statistically significant. Of course this finding does not imply that expected returns are constant; it only means that we cannot reject that hypothesis from a simple comparison of the first and second halves of the sample.

Several papers offer variations of two potential explanations for the declining performance of the value strategy: book value is worsening as a proxy for the value of assets in place (Choi, So, and Wang, 2021; Eisfeldt, Kim, and Papanikolaou, 2021; Goncalves and Leonard, 2021), or a series of shocks has widened the difference in multiples between growth and value stocks (Israel, Laursen, and Richardson, 2020; Arnott et al., 2021b). Each of these explanations imply time variation in expected returns for the value portfolio, meaning a methodology like ours is needed to address this possibility when making inferences about either near-term or long-term expected returns.

The investment, profitability, and size portfolios each show different patterns across subsamples. The investment portfolio's returns are largely consistent over time and statistically significant in both halves of the sample. The profitability portfolio follows the opposite pattern as the value portfolio, performing worse in the first half of the sample than the second, though again the difference is statistically insignificant. The size portfolio has small and statistically insignificant returns in both halves of the sample.[11]

2.3. Autocorrelation Estimates

As discussed in the Introduction, a number of recent studies document short-term autocorrelation in characteristic-sorted portfolio returns, or more generally analyze portfolio timing strategies that are inherently based on time-variation in expected returns. Before we estimate our model we present similar reduced form analyses for the four characteristic sorted portfolios that we analyze.

Estimates of the autocorrelation structure for individual portfolios are inherently imprecise because realized returns are much more volatile than plausible variations in expected returns. Furthermore, autocorrelation estimates have a well-document downward bias in small samples (Kendall, 1954; Marriott and Pope, 1954). To illustrate these difficulties, we

[11] We use value-weighted portfolios on both the long and short side, a construction which never produced a positive size effect.

first present autocorrelation estimates we generate in simulated samples from our model under a variety of assumptions about H and γ. In these simulations, we fix the model parameters and generate 50,000 samples with 226 observations, the number of quarters in our sample.[12] For each simulated return series, we estimate return autocorrelations using regressions of quarterly returns on averages of recent past returns:

$$r_t = a + b_L \left(\frac{1}{L} \sum_{l=1}^{L} r_{t-l} \right) + \epsilon_t, \tag{11}$$

where L is the number of past quarters that we average.

Panel A of Table 2 presents average values, as well as 95% confidence intervals, for the autocorrelation coefficient $\hat{b}$ across samples that are simulated under a variety of plausible parametric assumptions. Specifically, we include shocks with half-lives, H, that are 2.5, 5, and 10 years, and first-order autocorrelations, γ, that equal 2.5%, 5%, and 10%. For each parameterization, we also report the standard deviation of annualized Sharpe ratios conditional on past returns, σ(Cond. Sharpe), that is implied by that parameterization. These values indicate that even seemingly small values for γ (2.5% or 5%) can generate economically meaningful – though plausible – variations in conditional Sharpe ratios.

The first row of Panel A in Table 2 reports the coefficient estimates when $H = 0$ and thus returns are distributed i.i.d. Despite this, the average autocorrelation coefficients are negative due to the aforementioned downward bias. The reason for this bias is that we are estimating both the mean and autocorrelation of r_t. Because the estimate of the mean is always the in-sample average, the data will appear to mean-revert to this average even when there is no true mean reversion. The bias is stronger the longer the past return window because there is a smaller effective sample size. To take an extreme example, if we divide the

[12] The remaining model parameters we use for the simulations are $\mu = 0$ and $\sigma_r = 6.89\%$, the full-sample standard deviation for the value portfolio. These choices have no effect on our results as autocorrelation estimates with intercepts are invariant to linear transformations of r_t.

sample in two, one half will appear above-average and the other below-average, suggesting mean reversion, in every random sample.

The other rows in Panel A of Table 2 show the autocorrelation estimates that obtain when expected returns exhibit persistent variation, i.e., $H > 0$ and $\gamma > 0$. Two observations emerge from the reported estimates. First, the downward bias continues to affect autocorrelation estimates even when expected returns are persistent, especially with longer past return windows. Second, and more importantly, the confidence intervals for the autocorrelation estimates are quite wide and include negative values in every parameterization.

Panel B of Table 2 presents estimates of Equation (11) for the historical samples of the quarterly returns of the value, investment, profitability, and size portfolios. The first column shows that all four portfolios have positive b_4, indicating that past-year returns positively predict next-quarter returns. The second and third columns show that longer past-return windows have point estimates with varying signs.

We do not report standard errors or bias corrections in Panel B of Table 2 because Panel A shows that both depend heavily on the magnitude and persistence of the variations in expected returns. In most cases, the confidence intervals for different parameterizations in Panel A include all of the point estimates in Panel B, which means that one cannot reject any of the posited autocorrelation structures. Even economically large coefficients, such as $\hat{b}_{40} = -0.73$ for the investment portfolio, lie in the 95% confidence interval for everything from the i.i.d. parameterization to $H = 10$ and $\gamma = 5\%$.

For three out of the four characteristic-sorted portfolios the autocorrelation estimates do not reject any reasonable parameterization of our model. The size portfolio is the exception; one can convincingly reject the i.i.d. null at the one-year lag. In addition, the i.i.d. null hypothesis is rejected for one-year and ten-year returns with p-values of 0.0% and 4.7%, respectively, in a pooled regression that combines the returns of the four portfolios.[13] In a

[13] We conduct this test by estimating b_L in a pooled panel regression with portfolio fixed effects, and

pooled regression with only value, investment, and profitability, we reject the i.i.d null at one and ten year lags with p values of 0.1% and 15.2%, respectively. The reason we can jointly but not individually reject this null is that the coefficients for all four portfolios tend to be above the i.i.d. benchmark, but just not enough to reject on an individual basis.

3. OLS, GLS, and Maximum Likelihood Estimations

As the previous section shows, the autocorrelation patterns observed in the historical returns of characteristic-sorted portfolios are consistent with a range of different assumptions about the magnitude of persistent variation in conditional expected returns. In this section, we formally analyze the impact of such variation on the estimates of unconditional expected returns using OLS and GLS regressions and maximum likelihood estimations.

3.1. OLS and GLS Standard Error Corrections

We start by estimating unconditional expected returns μ and their accompanying standard errors using OLS regressions of the observed returns series r_t on a constant. The OLS estimate $\hat{\mu}^{OLS}$ is a consistent estimator of μ, even when conditional expected return vary as in our model. The correct standard errors for $\hat{\mu}^{OLS}$ depend on the covariance matrix of the residuals $\psi_t = r_t - \mu$.[14]

The typical approaches to adjust standard errors are the White (1980) correction for potential heteroskedasticity and Newey and West (1987), which corrects for both heteroskedasticity and autocorrelation up to a small number of lags. When expected returns are time-varying and persistent, both of the standard approaches produce understated standard errors. The reason is that persistent variations in expected returns generate small but long-lasting

comparing our estimate in observed data to the distribution of estimates in samples simulated by re-sampling observed months with replacement, retaining any cross-sectional relations among the portfolios.

[14]We use 'correct standard errors' as an informal shorthand for the correct specification of the asymptotic distribution of $\hat{\mu}$.

correlations in ψ_t that extend beyond the windows considered by Newey and West (1987). Formally, our model generates the residuals

$$\psi_t = \mu_{t-1} - \mu + \epsilon_t, \tag{12}$$

which implies that ψ_t have long-lasting autocorrelations due to persistent variations in μ_t. As we show using simulations in Section 3.2, even if we extend the number of lags in Newey and West (1987) to match or exceed the half-lives of shocks to expected returns, standard errors remain under-estimated because the lags are too large a fraction of the observed data for the asymptotic results in Newey and West (1987) to hold.

If we know the values of H and γ, we can correct OLS standard errors for the resulting autocorrelation in ψ_t using the structure of our model. When residuals have a known covariance matrix $\mathbf{\Sigma}$, the asymptotic variance of $\hat{\mu}^{OLS}$ is:

$$T^2 \text{Var}(\hat{\mu}^{OLS}) = \mathbf{1}'\mathbf{\Sigma}\mathbf{1}, \tag{13}$$

where $\mathbf{1}$ is a $T \times 1$ vector of ones (see Section 4.5 of Cameron and Trivedi (2005)). In our model, $\mathbf{\Sigma}$ is fully specified by H and γ, as shown in Equation (4).

We also estimate μ in a GLS regression, which uses the covariance matrix Σ implied by H and γ to adjust both the standard error and the point estimate $\hat{\mu}$. While the OLS estimates are based on an equally-weighted average of the returns in the sample, the GLS estimates utilize an average weighted by the amount of orthogonal information each observation contains about the unconditional expected return. When conditional expected returns exhibit persistent variations, the observations in the middle of the sample are somewhat redundant because they 'over-sample' the same epoch of conditional expected returns. As illustrated by the first panel of Figure 1, GLS therefore overweights observations at the beginning and end of the sample. This effect is larger for higher H, and reverses when γ is negative.

3.2. OLS and GLS Estimation Results

Table 3 shows how different assumptions about the parameters H and γ affect OLS and GLS point estimates and standard errors for unconditional expected returns μ. As a point of comparison, we also provide the Newey-West standard errors calculated with lags equal to the half-life H (expressed as the number of quarters). Across all four portfolios, Newey-West standard errors increase little or not at all with H. Model-implied standard errors, by contrast, more than double when $H = 5$ years and $\gamma = 5\%$ – a scenario that Table 2 shows we cannot reject for any of these portfolios.[15] With this autocorrelation structure, we no longer reject the null that $\mu = 0$ at the 5% level for any of the four portfolios.[16]

The GLS correction to $\hat{\mu}$ discussed above has minimal impact on the point estimates for the profitability portfolio, but noticeably decreases point estimates for value, investment, and size. The reason for this decrease is that the latter three portfolios had unusually low returns at the beginning and/or ends of the sample, which GLS infers as being more independent from the observations in the middle.

The overall takeaway from Table 3 is that accounting for plausible degrees of persistent variation in conditional expected returns results in substantially higher standard errors for the estimates of unconditional expected returns. Thus, the standard assumption in the literature of i.i.d. returns tends to overstate the precision with which the expected returns of characteristic-sorted portfolios are estimated.

[15] The only exception is that the size portfolio has a higher b_4 than predicted by this scenario, meaning it is even more persistent and standard errors for $\hat{\mu}$ should be even larger.

[16] In untabulated results, we simulate the small-sample distribution of average returns using our model for each H and γ. We find that the asymptotic model-implied standard errors are almost identical to their small-sample counterparts, meaning that in this setting the bias in Newey and West (1987) is due to the mis-specification of autocorrelation rather than sample size.

3.3. *Maximum Likelihood Tests*

Our next set of tests use maximum likelihood estimates of our statistical model. Unlike the least squares regressions in the previous subsections, maximum likelihood estimation requires distributional assumptions, but allows us to estimate all of the model parameters, including those that govern the persistent variation of returns.

3.3.1. *Tests for $\mu = 0$*

We test the null hypothesis that the unconditional expected return μ equals zero under a variety of assumptions about the structure of time-variations in conditional expected returns. For each assumption, we estimate the model using maximum likelihood twice, first with no constraints on μ and then restricting $\mu = 0$. In both cases we restrict H to be less than or equal to 20 years because μ is not identified when H approaches infinity. Using these estimates, we compute the p-value for the $\mu = 0$ null using a likelihood ratio test.

The first set of columns in Panel A of Table 4 test whether $\mu = 0$ assuming no time variation in expected returns (i.e. $H = 0$ so returns are i.i.d.). Testing $\mu = 0$ under this assumption is analogous to using OLS with no standard error correction, and so the likelihood-ratio p-values are quite low for value, investment, and profitability.

The next set of columns in Panel A of Table 4 relax the i.i.d. assumption and estimate the values of H and γ that maximize the likelihood of observing the historical data.[17] We find that the p-value for the hypothesis that $\mu = 0$ increases for all three portfolios once we allow for the possibility of time-varying means. As the panel shows, when μ is restricted to be zero, the model fits the data by using larger values of H combined with positive values of γ. This combination explains the observed past returns as arising from positive realizations

[17] As discussed in Section 1, we estimate the parameters summarizing the covariance matrix of returns $\theta = [\mu, H, \sigma_r, \gamma]$ directly rather than the underlying parameters $\Omega = [\mu, \lambda, \sigma_\epsilon, \sigma_\delta, \rho,]$ because the later are not fully identified. We restrict our estimates of θ to the range for which there exists at least one possible Ω yielding the same covariance matrix of returns.

of *persistent* expected returns that eventually dissipate. Because such persistent variation is difficult to reject empirically, it offers a plausible alternative to large μ, increasing the p-values for rejecting the $\mu = 0$ hypothesis above 5% in all cases.

3.3.2. Restrictions on H and γ

Next, we use maximum likelihood estimates to assess the plausibility of the various assumptions about H and γ that we made in our linear regression analyses in Section 3. For each assumption, we re-estimate the model by restricting H and γ to their assumed values, and use likelihood ratios relative to the unrestricted model to test whether we can reject the restriction. We also use likelihood ratios to test whether we can reject the hypothesis that $\mu = 0$ given the restrictions on H and γ.

Panel B of Table 4 presents the results. For value, investment, and profitability, the likelihood ratio tests cannot reject *any* of the restrictions on the structure of time-variation in expected returns, including the IID hypothesis and the possibility of extremely-persistent and economically large shocks ($H = 10$ years, $\gamma = 5\%$). Consistent with the evidence in Table 3, Panel B of Table 4 also shows that these alternatives have material affects on our inferences about $\mu = 0$, with p-values for $\mu = 0$ going from below 2% in the i.i.d. case to above 10% in some specifications for all three characteristic-sorted portfolios.

While the results for value, investment, and profitability suggest that the data offer very little guidance about the magnitude of persistent variations in expected returns, the results for the size portfolio show that this is not always the case. For size, as the reduced-form evidence in Table 2 suggests, there is strong autocorrelation in returns with a half life between one and two years. As a result, Panel B of Table 4 shows the maximum likelihood estimator strongly rejects both the i.i.d. restriction and the restriction that $H = 10$. This contrast shows it is not always true that 'anything goes' when our model is applied to data, but instead that the returns for value, investment, and profitability are particularly inconclusive.

4. Bayesian Analysis

The analysis in the previous section shows that the historical data are consistent with a variety of substantially different return generating processes. A natural next step is to ask how investors with different prior views about these alternatives interpret the historical data in making their investment decisions. We examine this normative question in this section. Specifically, we pursue a Bayesian analysis that specifies prior likelihoods of different model parameterizations and uses the observed data to calculate posterior likelihoods. We also compute posteriors for moments that are likely to affect investors' portfolio decisions, such as the near-term and long-term Sharpe ratios for each portfolio.

4.1. Prior Beliefs

We specify prior beliefs over a transformation of the parameters that is focused on Sharpe ratios rather than expected returns because Sharpe ratios have a clearer connection to asset pricing theory and economic intuition.[18] Specifically, we specify priors over μ_{sr}, the unconditional Sharpe ratio of portfolio returns; H, the half-life of shocks to expected returns; σ_r, the unconditional standard deviation of portfolio returns; σ_{sr}, the standard deviation of conditional Sharpe ratios; and ρ, the correlation between unexpected returns shocks to expected returns. We tabulate μ_{sr}, H, σ_r, and σ_{sr} in annualized units.

These parameters map into the underlying model parameters Ω as follows:

$$\mu_{sr} = \frac{\mu}{\sigma_\epsilon}, \qquad H = \frac{\log(0.5)}{\log(\lambda)}N, \qquad \sigma_r^2 = \frac{\lambda^2\sigma_\delta^2}{1-\lambda^2} + \sigma_\epsilon^2,$$

$$\sigma_{sr}^2 = \frac{\lambda^2\sigma_\delta^2/(1-\lambda^2)}{\sigma_\epsilon^2}, \qquad \rho = \rho.$$

[18]Another advantage of this transformation is that the same Sharpe ratio priors can be applied across all characteristic-sorted portfolios and regardless of the amount of financial leverage applied to a portfolio investment.

We consider a variety of priors on μ_{sr} and H, summarized in Panel A of Table 5. For μ_{sr}, we first consider normal prior distributions centered at -0.4, 0, and 0.4, all with a standard deviation of 0.25. These priors can be interpreted as beliefs that the portfolio has unconditional Sharpe ratios that are likely to be negative, likely to be near-zero, or likely to be positive.[19] For the value portfolio, these three analysts can be viewed as having growth, neutral, or value inclinations. We also examine an uninformative prior where any μ_{sr} between -2 and 2 is equally likely.[20]

To illustrate the effect that prior beliefs about H have on Bayesian inferences about expected returns, we consider three dogmatic priors and one agnostic prior. The dogmatic priors assert that $H = 0$ (making returns i.i.d.), $H = 2.5$ years, or $H = 5$ years with certainty. The agnostic prior, by contrast, views H as unknown and uniformly distributed between 0 and 10 years.

We consider uniform priors over wide ranges for the remaining parameters. The prior for the volatility of annual returns, σ_r, is uniformly distributed between 10% and 20%.[21] The prior for the standard deviation of conditional Sharpe ratios, σ_{sr}, is uniformly distributed between 0 to 1.[22] The correlation between unexpected returns and shocks to expected returns, ρ, is likely to be negative given the inverse relation between prices and expected returns. In contrast to market returns, however, characteristic-sorted portfolio returns should be driven primarily by cash-flow news rather than discount rate news. Based on these observations, we specify the prior on ρ to be uniformly distributed in the interval $[-0.5, 0]$.

[19] We use ± 0.4 as the center for our Sharpe ratio distributions to roughly match the US equity market's estimated Sharpe ratio.

[20] Uniform distributions over wider supports give nearly-identical results because the data strongly reject unconditional Sharpe ratios above 2 or below -2 for for the portfolios we study when H is less than 10 years.

[21] This is the only prior that would need to shift as the leverage applied to a portfolio changes.

[22] In the most extreme case, this implies that conditional Sharpe ratios occasionally deviate from their unconditional means by as much as 2. While this extreme case is implausible given active asset management chasing and reducing large Sharpe ratio opportunities, as we emphasize below the true conditional Sharpe ratio may not be observable, and Sharpe ratios that are achievable by conditioning only on past returns will be much less extreme.

To provide further intuition for the set of priors we study, Panel B of Table 5 presents the distribution of economically-intuitive moments implied by the priors specified in Panel A. We calculate these moments by simulating 50,000 draws from each prior and calculating the value of each moment implied by each parameter draw. The first set of columns shows that μ, the unconditional expected return, has a prior mean that is about -5.3%, zero, or 5.3%, depending on the prior specification. The middle set of columns show that the prior mean values of one-lag return autocorrelation, γ, are positive in all cases, indicating that the positive effect due to persistence of expected returns outweighs the negative effect due to $\rho < 0$, though the 95% confidence intervals do include negative values.

The parameter σ_{sr} governs the volatility of the true conditional Sharpe ratio of returns, which is not observable to econometricians and may not be perfectly observable to investors in practice. If investors instead have to use past returns, combined with beliefs about model parameters, to forecast future returns, the relevant economic magnitude is the volatility of Sharpe ratios conditional on past returns. We compute this volatility in the steady state, conditioning on an infinite history of past returns to forecast next-period returns. These forecasts remain imperfect in the steady state because the current μ_t is unobservable and can only by imperfectly filtered from past return observations.

The last set of columns in Panel B in Table 5 show that these conditional Sharpe ratios can vary substantially as the priors about H increase and allow for more persistence in returns. However, despite the relatively wide range of priors for σ_{sr}, the ability to predict portfolio returns using past returns is somewhat limited, with the prior mean of $\sigma(\textit{Cond. Sharpe})$ around half the prior mean of σ_{sr}.

4.2. Posteriors for Unconditional Expected Returns and Sharpe Ratios

We estimate the posterior distribution for each of 16 priors (four for H and μ_{sr} each) and each of four characteristic-based portfolios, making 64 prior-data pairs. We generate

50,000 posterior draws for each prior-data pair using the algorithm detailed in Appendix B. We compute the posterior for the full parameter vector Ω. However, as discussed above, we cannot separately identify all five parameters in Ω in our frequentist estimations, which manifests in extremely wide posteriors for ρ and σ_δ in the Bayesian analysis. For each posterior we therefore compute the four identified components of θ, namely μ, σ_r, H, and γ. Also, instead of γ, we present σ(Cond. Sharpe), which is easier to interpret.

Tables 6 through 9 and Figures 2 through 5 report the results for each of the four characteristic-sorted portfolios. As these results show, different priors result in substantially different inferences about unconditional expected returns μ and Sharpe Ratios μ_{sr}. For instance, Bayesian analysts' posteriors for μ and μ_{sr} are wider as their priors for H increase, making the possibility of zero or negative μ much more likely. The intuition is the same as for the frequentist analysis above: the data do not strongly reject the possibility that unconditional expected returns are zero or negative and that the historical performance is explained by persistent but dissipating positive shocks to conditional expected returns.

The Bayesian analysis also produces an insight that is distinct from the frequentist analysis: the extent to which priors about μ_{sr} affect posteriors about μ and μ_{sr} depends on the analyst's prior about H. Bayesian analysts who believe $H = 0$ largely agree about μ and μ_{sr} despite large differences in priors. On the other hand, Bayesian analysts whose priors are that H is, or might be, large have substantial differences in their posterior beliefs about μ and μ_{sr} despite observing 56 years of data. For example, means of posteriors about the value portfolio's μ_{sr} are clustered between 0.25 and 0.32 for the $H = 0$ prior, but vary from 0.14 to 0.30 for the $H \sim U(0, 10)$ prior.

The reason priors about H matter for posteriors about μ is that Bayesian analysts 'shrink' observed in-sample averages towards the mean of their prior, and the extent of this adjustment depends on H. If the analyst believes H equals zero and thus returns are i.i.d., the data are more informative about unconditional expected returns and thus the posterior hews closer to the in-sample average. If the analyst believes H is or may be larger than zero, the data are less informative and so the posterior depends more on their prior. As a result, the posterior Sharpe ratios vary more across priors for μ_{sr} (rows in Figures 2 through 5) when priors for H are larger.

Tables 6 through 9 also show that the posterior distributions for H and σ(Cond. Sharpe) differ very little from the corresponding prior distributions for the value, investment, and profitability portfolios, which is consistent with the evidence in Tables 2 and 4 that the data offer little guidance on the autocorrelation structure for these portfolios. The size portfolio, by contrast, has posterior distributions for H that are farther from zero than the prior and have means below the prior mean of five years, indicating the data push Bayesian analysts towards lower H.[23] Furthermore, the mean posterior σ(Cond. Sharpe) is much higher for size than the other portfolios, and unlike the other portfolios its confidence interval excludes zero, meaning the evidence tilts in favor of positive autocorrelation. Still, posteriors for the size portfolio's H and σ(Cond. Sharpe) are quite wide, leaving room for many potential interpretations of the data.

Just as priors about H affect posteriors about μ_{sr}, priors about μ_{sr} also affect posteriors about H. Analysts with bearish priors on μ_{sr} are more amenable to interpreting the observed positive Sharpe ratios as arising from extremely persistent (high H) variations in conditional Sharpe Ratios, leading them to tilt their H posteriors higher than analysts with positive priors on μ_{sr} that are more comfortable with large unconditional Sharpe Ratios explaining the observed data. Tables 6 and 9 show this is indeed the case: for the portfolios with substantial average returns in our sample (value, investment, and profitability), posteriors about H have higher means for analysts with bearish priors about μ_{sr}. This effect is small quantitatively, however, because there is so little power in measuring H.

[23]The prior for H is uniformly distributed between 0 and 10, meaning the 95% confidence interval is $[0.25, 9.75]$.

4.3. Posteriors for Conditional Sharpe Ratios

In addition to forming inferences about unconditional expected returns and Sharpe ratios, the Bayesian analysis of our model allows us to compute posterior distributions of conditional Sharpe ratios through each period in our historical sample. These conditional Sharpe ratio distributions are computed using the posterior distributions for model parameters based on the full sample of data, and conditioning on the full sample of returns.[24] Figure 6 plots the time-series of posterior means for the four characteristic-sorted portfolios for the 'agnostic' prior specification with $\mu_{sr} \sim U(-2, 2)$ and $H \sim U(0, 10)$. As the figure shows, the variations in conditional Sharpe ratios are economically substantial but plausible, generally varying between 0 and 0.8 on an annualized basis. The exception is size, which, despite having estimates of unconditional Sharpe ratios near zero, has conditional Sharpe ratios varying from -0.75 to above one. These fluctuations arise due to a combination of the large variations in rolling average returns presented in Figure **??** and the posterior beliefs tilted towards strong persistence in Table 9.

Finally, we compute posterior distributions of forward-looking conditional Sharpe ratios for the quarter immediately following the end of our sample period (Q1 of 2020). These posteriors differ from the unconditional posteriors because they use nearby trends in portfolio returns to extrapolate to future performance. When $\gamma > 0$, the extrapolation is positive, meaning that near-term expected returns are higher (lower) than long-term expected returns when recent returns are higher (lower) than the full-sample average. When $\gamma < 0$, the extrapolation is negative, meaning near-term expectations reverse the recent trends. Figure 1 illustrate these patterns by plotting influence functions, defined as the impact each ob-

[24]Note that even with the full sample of data we do not observe conditional expected returns in each period and instead need to estimate them using realized returns and model parameters. These results should be interpreted as a measure of full-sample economic significance (such as an R^2) rather than feasible Sharpe ratios that inform real-time trading strategies.

servation has on the posterior belief about conditional expected returns in 2020 for a few pararametrizations. The Bayesian posteriors integrate these influence functions across the posterior distribution of parameters.

Panel B of Figures 2 through 5 present the mean and 95% confidence intervals for posterior beliefs about the 2020 conditional Sharpe ratios of each portfolio. As the first rows shows, conditional and unconditional Sharpe ratios are the same when $H = 0$. When $H > 0$, however, both value and investment have smaller conditional Sharpe ratios than unconditional. [25] For the value portfolio, which had particularly poor recent performance as discussed above, the pessimistic or neutral Bayesian analysts believes conditional Sharpe ratios in 2020 are centered near zero and could even be quite negative.

Because the profitability portfolio performed better in recent years than earlier in the sample, we find the opposite effect in Figure 4: posteriors about 2020 Sharpe ratios are *higher* than posteriors about the unconditional Sharpe ratio. As with the bearish view of value and investment, the bullish view for profitability is stronger for larger values of H than smaller ones.

5. Conclusion

Time-variation in expected returns is an overlooked but important source of uncertainty about the long-term performance of characteristic-based portfolios. The uncertainty is both direct – the data offer little guidance on the magnitude and persistence of expected return shocks – and indirect – the fact that there may be persistent variations in conditional expected returns causes the data to be much less informative about unconditional expected returns. Because the data are less informative, analysts with different priors can have substantially different posterior views despite observing more than five decades of return data.

[25] These effects are stronger for the 2020 forecasts than they are for longer-horizon forecasts because the model predicts that the temporary conditional expected returns that prevail today will mean revert over time.

Our results have implications for practitioners and academics interested in characteristic-sorted portfolios. Financial institutions now offer a multitude of relatively passive investment products, such as ETFs and mutual funds, that aim to exploit the long-term links between returns and characteristics that have been identified in academic research. At the same time, there exist active hedge funds that attempt to time the variations in characteristic return premia that are described in the more recent literature. Our Bayesian analysis offers guidance for both groups as to how much they should use recent performance to guide their decisions, as well as the near- and medium-term performance outlooks for these portfolios.

Overlooking the uncertainty caused by time-variation in expected returns could also offer a partial explanation for the ever-increasing number of characteristics or factors that predict returns (McLean and Pontiff (2016), Harvey, Liu, and Zhu (2016)). The ease with which finance research uncovers such anomalies raises doubts about whether the documented return patterns point to genuinely long-term return premia. We offer a potential alternative explanation: historical returns are often driven by economic or behavioral conditions that may persist for decades but not indefinitely.

Appendix A. Autocorrelations in Monthly Returns

As described in Section 2.1, we analyze a sample of quarterly returns for characteristic-based portfolios rather than the monthly returns typically studied in the literature. We do so because monthly returns exhibit strong first-order autocorrelations that may be caused by lead-lag effects, under-reaction, or some other transitory source of persistence. While an interesting topic on its own, this form of autocorrelation is not the topic of this paper, namely slow-moving but persistent variations in expected returns. These two phenomenon are likely to co-exist, and so to avoid biasing our estimates towards large but quickly-reverting variations, we use quarterly rather than monthly data.

Appendix Figure 1 illustrates the magnitude of the autocorrelations at monthly lags $l = 1$ through $l = 60$ for the four portfolios we study, as well as the autocorrelations estimated in a pooled regression including all four portfolios. The first-order autocorrelation is the single largest coefficient for any of the 60 months for value, investment, and profitability, as well as in pooled regressions. Quarterly data does not exhibit a strong first-order autocorrelation (see Table 2) because the two- and three-month autocorrelation in Appendix Figure 1 are much smaller and statistically insignificant.

Appendix B. Sampling Bayesian Posteriors

We draw samples of $N = 50,000$ observations from the posterior distribution of model parameters Ω^{post} using the following procedure:

1. Draw N observations Ω_i^{prior}, $i \in [1, N]$ from the prior distribution.

2. Accept Ω_1^{prior} as the first observation of the posterior distribution $\Omega_1^{\text{posterior}}$.

3. For observations $i = 2 \ldots N$:

   (a) Evaluate the conditional likelihood of the data D given the ith draw from the prior parameters as well as the $i-1$st draw from the posterior parameters:

$$\mathcal{L}^{\text{propose}} = \mathcal{L}(D|\Omega = \Omega_i^{\text{prior}}), \tag{14}$$

$$\mathcal{L}^{\text{previous}} = \mathcal{L}(D|\Omega = \Omega_{i-1}^{\text{posterior}}). \tag{15}$$

   (b) If $\mathcal{L}^{\text{propose}} \geq \mathcal{L}^{\text{previous}}$, accept $\Omega_i^{\text{posterior}} = \Omega_i^{\text{prior}}$.

   (c) If $\mathcal{L}^{\text{propose}} \leq \mathcal{L}^{\text{previous}}$, accept $\Omega_i^{\text{posterior}} = \Omega_i^{\text{prior}}$ with probability $\frac{\mathcal{L}^{\text{propose}}}{\mathcal{L}^{\text{previous}}}$, and otherwise retain $\Omega_i^{\text{posterior}} = \Omega_{i-1}^{\text{prior}}$.

Appendix Figure 1: Monthly Autocorrelograms

This figure presents autocorrelations of monthly returns for value, investment, profitability, and size portfolios, as defined in the header of Table 1, as well the autocorrelation estimated in a pooled regression containing all four portfolios. We estimate the autocorrelation for each lag l independently. The horizontal lines represent the 95% confidence interval for autocorrelation coefficients under the zero-autocorrelation null hypothesis. Our sample consists of 678 monthly observations from Q3 1963 through Q4 2019.

Panel A: Value

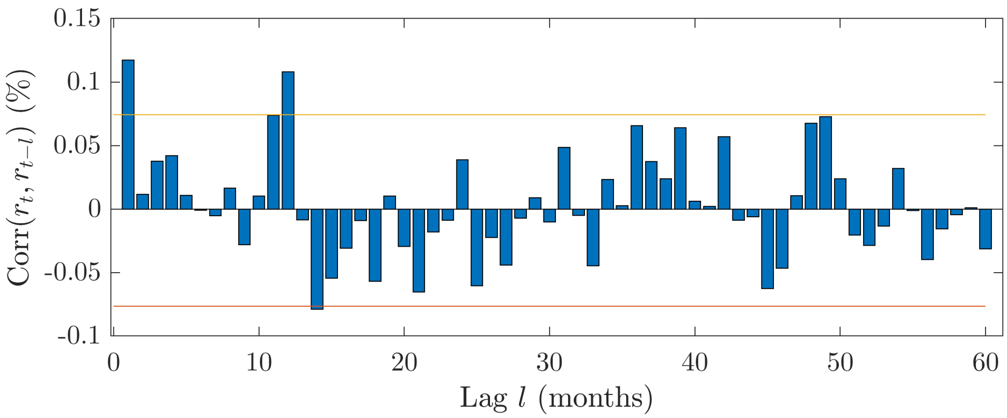

Panel B: Investment

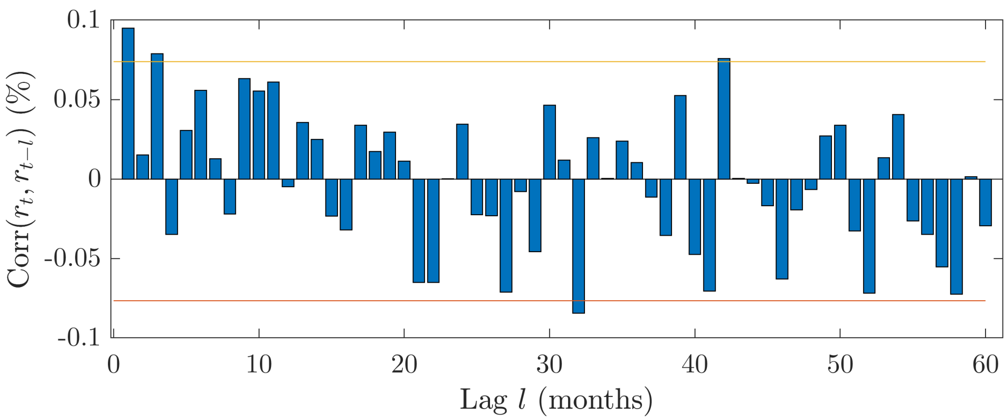

Appendix Figure 1: Monthly Autocorrelograms (continued)

Panel C: Profitability

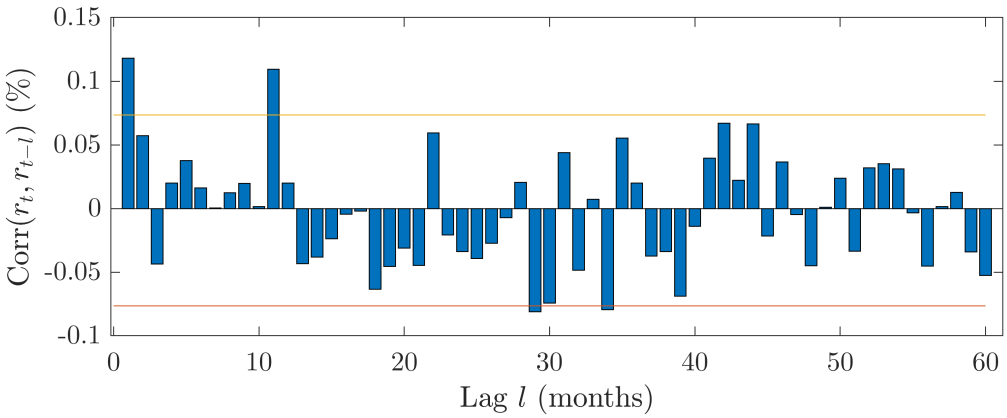

Panel D: Size

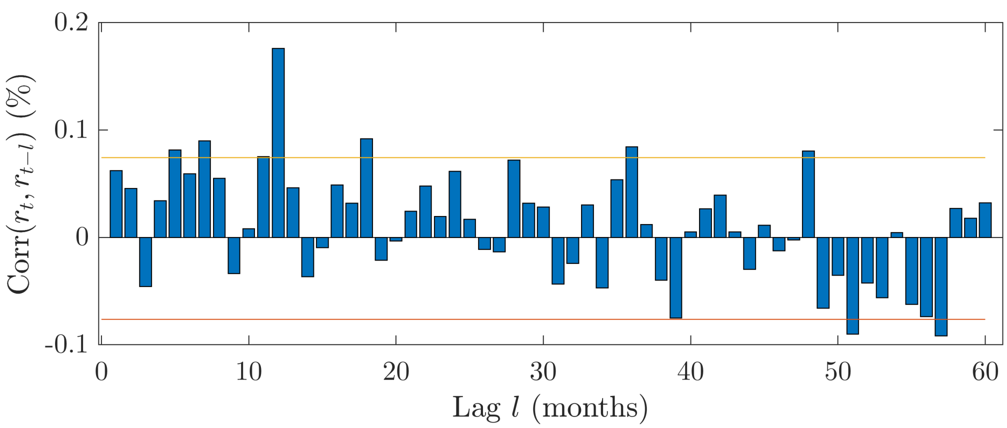

Panel E: Pooled

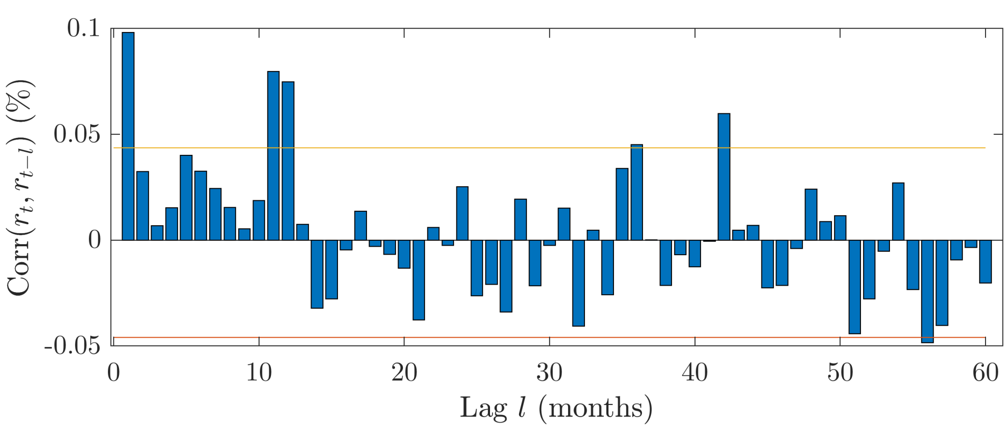

References

Altı, Aydoğan and Sheridan Titman, 2019, A dynamic model of characteristic-based return predictability, *The Journal of Finance* 74, 3187–3216.

Arnott, Robert D, Mark Clements, Vitali Kalesnik, and Juhani T Linnainmaa, 2021a, Factor momentum, *Available at SSRN 3116974* .

Arnott, Robert D, Campbell R Harvey, Vitali Kalesnik, and Juhani T Linnainmaa, 2021b, Reports of values death may be greatly exaggerated, *Financial Analysts Journal* 77, 44–67.

Asness, Cliff, John Liew, Lasse Heje Pedersen, and Ashwin Thapar, 2021, Deep value, *The Journal of Portfolio Management* 47, 11–40.

Avramov, Doron, Si Cheng, Amnon Schreiber, and Koby Shemer, 2017, Scaling up market anomalies, *The Journal of Investing* 26, 89–105.

Barberis, Nicholas, 2000, Investing for the long run when returns are predictable, *The Journal of Finance* 55, 225–264.

Cameron, A Colin and Pravin K Trivedi, *Microeconometrics: methods and applications* (Cambridge University Press 2005).

Choi, Ki-Soon, Eric C So, and Charles CY Wang, 2021, Going by the book: Valuation ratios and stock returns, *Available at SSRN 3854022* .

Ehsani, Sina and Juhani T Linnainmaa, 2021, Factor momentum and the momentum factor, *The Journal of Finance* Forthcoming.

Eisfeldt, Andrea L, Edward Kim, and Dimitris Papanikolaou, 2021, Intangible value, *Available at SSRN 3720983* .

Fama, Eugene F and Kenneth R French, 2015, A five-factor asset pricing model, *Journal of Financial Economics* 116, 1–22.

Fama, Eugene F and Kenneth R French, 2021, The value premium, *The Review of Asset Pricing Studies* 11, 105–121.

Ferson, Wayne E, Sergei Sarkissian, and Timothy T Simin, 2003, Spurious regressions in financial economics?, *The Journal of Finance* 58, 1393–1413.

Gârleanu, Nicolae, Leonid Kogan, and Stavros Panageas, 2012, Displacement risk and asset returns, *Journal of Financial Economics* 105, 491–510.

Goncalves, Andrei and Gregory Leonard, 2021, The fundamental-to-market ratio and the value premium decline, *Available at SSRN 3573444* .

Greene, William H, *Econometric Analysis,* 4^{th} *Edition* (Prentice Hall 2000).

Gupta, Tarun and Bryan Kelly, 2019, Factor momentum everywhere, *Journal of Portfolio Management* 45, 13–36.

Harvey, Campbell R, Yan Liu, and Heqing Zhu, 2016, and the cross-section of expected returns, *The Review of Financial Studies* 29, 5–68.

Israel, Ronen, Kristoffer Laursen, and Scott Richardson, 2020, Is (systematic) value investing dead?, *The Journal of Portfolio Management* 47, 38–62.

Johannes, Michael, Arthur Korteweg, and Nicholas Polson, 2014, Sequential learning, predictability, and optimal portfolio returns, *The Journal of Finance* 69, 611–644.

Kandel, Shmuel and Robert F Stambaugh, 1996, On the predictability of stock returns: an asset-allocation perspective, *The Journal of Finance* 51, 385–424.

Kendall, Maurice G, 1954, Note on bias in the estimation of autocorrelation, *Biometrika* 41, 403–404.

Kogan, Leonid, Dimitris Papanikolaou, and Noah Stoffman, 2020, Left behind: Creative destruction, inequality, and the stock market, *Journal of Political Economy* 128, 855–906.

Lewellen, Jonathan, 2002, Momentum and autocorrelation in stock returns, *The Review of Financial Studies* 15, 533–564.

Linnainmaa, Juhani T and Michael R Roberts, 2018, The history of the cross-section of stock returns, *Review of Financial Studies* 31, 2606–2649.

Marriott, FHC and JA Pope, 1954, Bias in the estimation of autocorrelations, *Biometrika* 41, 390–402.

McLean, R David and Jeffrey Pontiff, 2016, Does academic research destroy stock return predictability?, *The Journal of Finance* 71, 5–32.

Newey, Whitney K and Kenneth D West, 1987, A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix, *Econometrica* 55, 703–708.

Novy-Marx, Robert, 2013, The other side of value: The gross profitability premium, *Journal of Financial Economics* 108, 1–28.

Pástor, L'uboš, 2000, Portfolio selection and asset pricing models, *The Journal of Finance* 55, 179–223.

Pástor, L'uboš and Robert F Stambaugh, 2009, Predictive systems: Living with imperfect predictors, *The Journal of Finance* 64, 1583–1628.

Pástor, L'uboš and Robert F Stambaugh, 2012, Are stocks really less volatile in the long run?, *The Journal of Finance* 67, 431–478.

Shiller, Robert J, *Irrational exuberance* (Princeton university press 2000).

Wachter, Jessica A and Missaka Warusawitharana, 2009, Predictable returns and asset allocation: Should a skeptical investor time the market?, *Journal of Econometrics* 148, 162–178.

White, Halbert, 1980, A heteroskedasticity-consistent covariance matrix estimator and a direct test for heteroskedasticity, *Econometrica* 817–838.

Figure 1: GLS Influence Functions and Return Autocorrelations

This figure presents the weights GLS uses when calculating unconditional (Panel A) and conditional end-of-sample (Panel B) average returns under a variety of assumptions about the correlation matrix of returns. H represents the half life of shocks to expected returns, in years. γ is the first-order autocorrelation of returns driven by persistence shocks to expected returns and correlations between realized returns and shocks to expected returns.

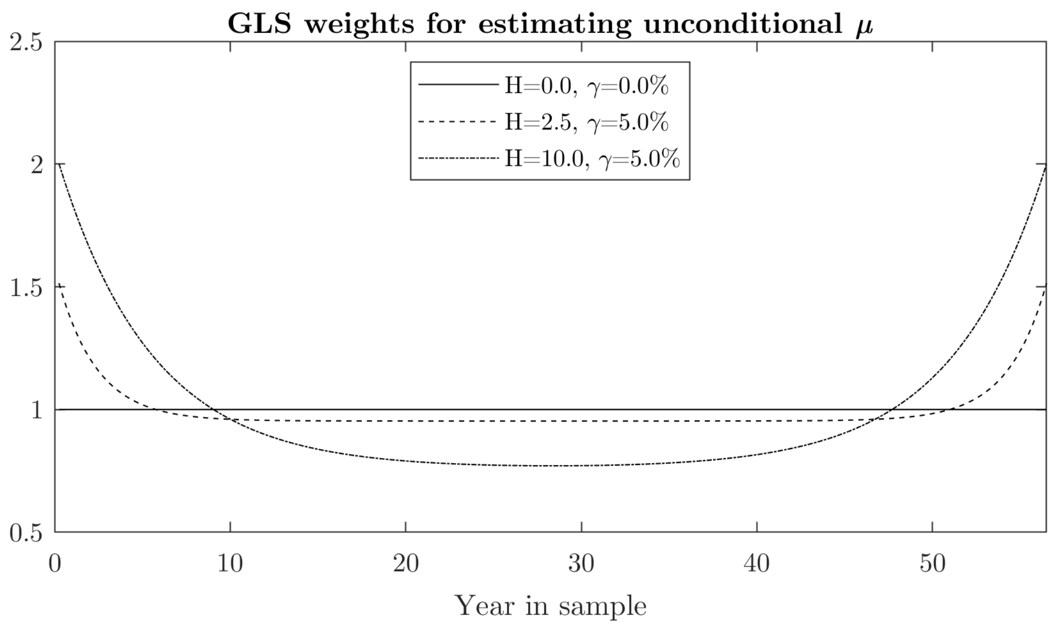

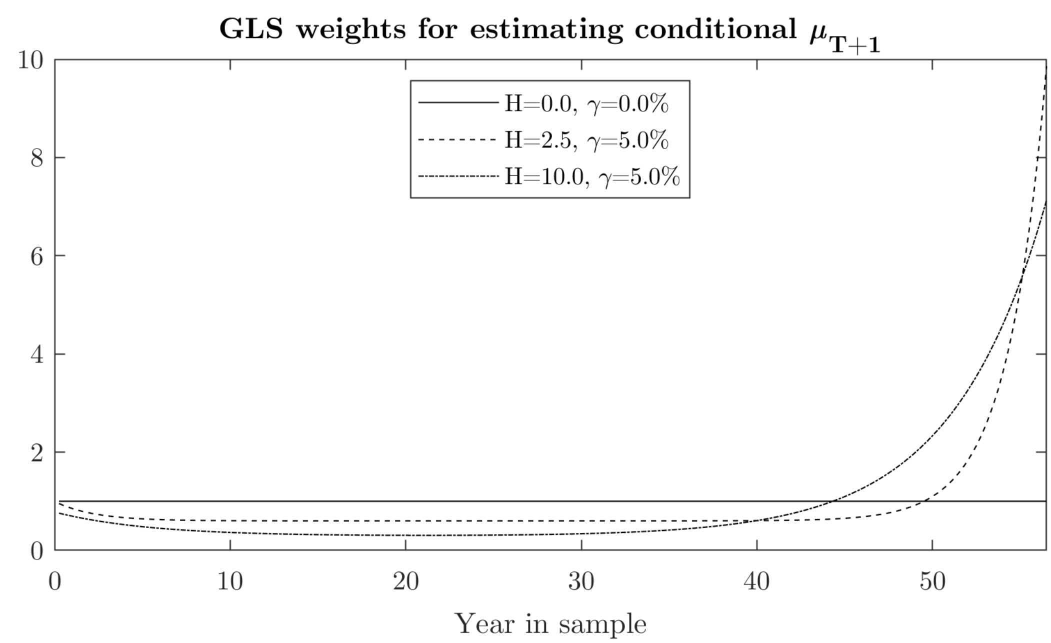

Figure 2: Posterior Sharpe Ratios for Value Portfolio

This figure presents posterior beliefs for the Sharpe Ratios of a portfolio formed by sorting stocks on value, as defined in the header of Table 1. Panel A presents beliefs about the unconditional Sharpe Ratio, Panels B and C about the conditional Sharpe Ratios in Q1 of 2020 and Q4 of 2029, respectively. Each panel shows posteriors for twelve possible priors, combining three possible priors for the unconditional Sharpe Ratio, μ_{sr}, with four possible priors for the half-life of shocks to conditional Sharpe Ratio, H. All Sharpe Ratios and H are annualized. Our sample consists of 226 quarterly observations from Q3 1963 through Q4 2019.

Panel A: Unconditional Sharpe Ratio

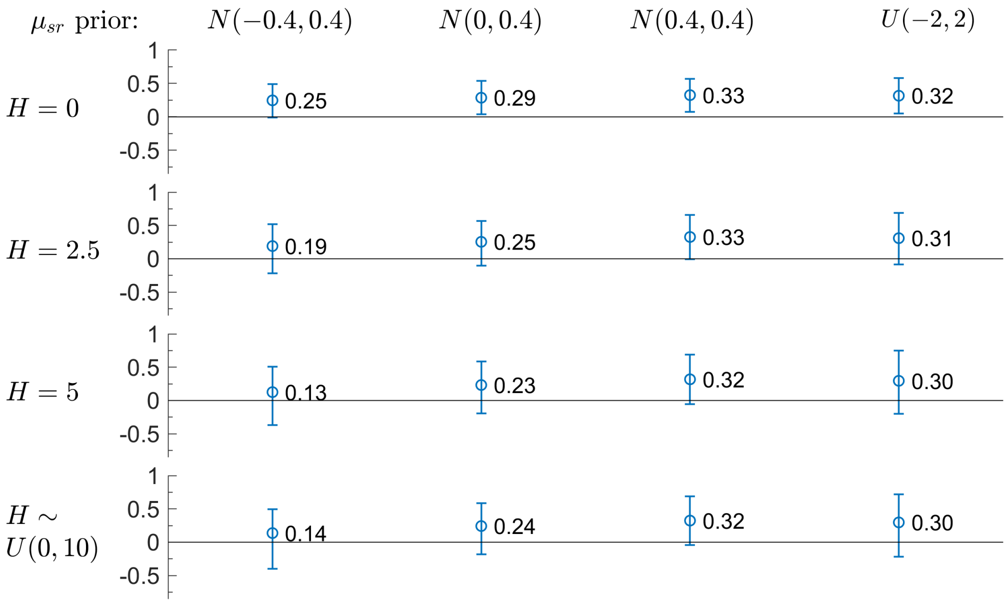

Panel B: Sharpe Ratio for 2020 Conditional on 1963–2019 Returns

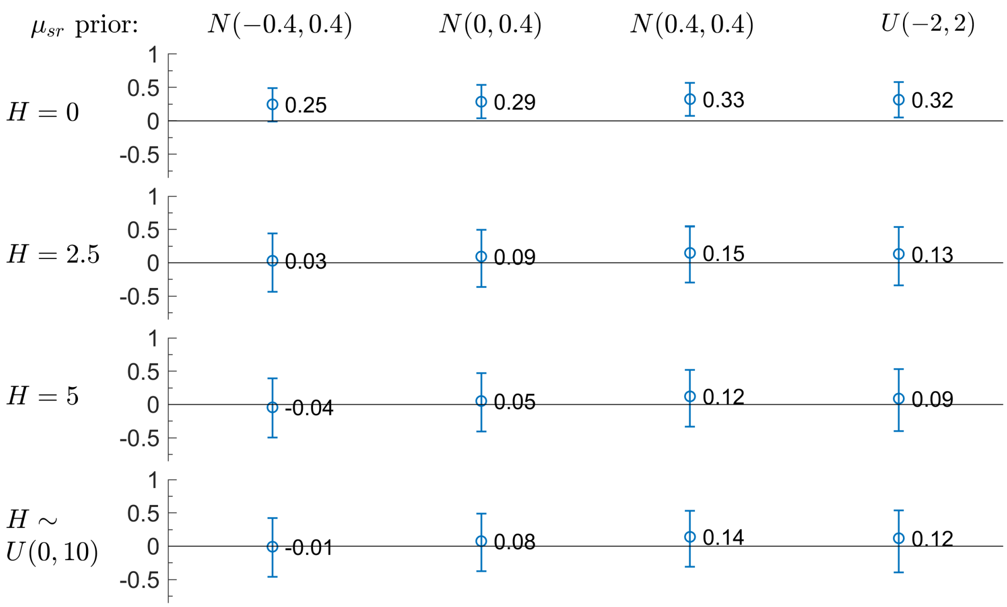

Figure 3: Posterior Sharpe Ratios for Investment Portfolio

This figure presents posterior beliefs for the Sharpe Ratios of a portfolio formed by sorting stocks on investment, as defined in the header of Table 1. Panel A presents beliefs about the unconditional Sharpe Ratio, Panels B and C about the conditional Sharpe Ratios in Q1 of 2020 and Q4 of 2029, respectively. Each panel shows posteriors for twelve possible priors, combining three possible priors for the unconditional Sharpe Ratio, μ_{sr}, with four possible priors for the half-life of shocks to conditional Sharpe Ratio, H. All Sharpe Ratios and H are annualized. Our sample consists of 226 quarterly observations from Q3 1963 through Q4 2019.

Panel A: Unconditional Sharpe Ratio

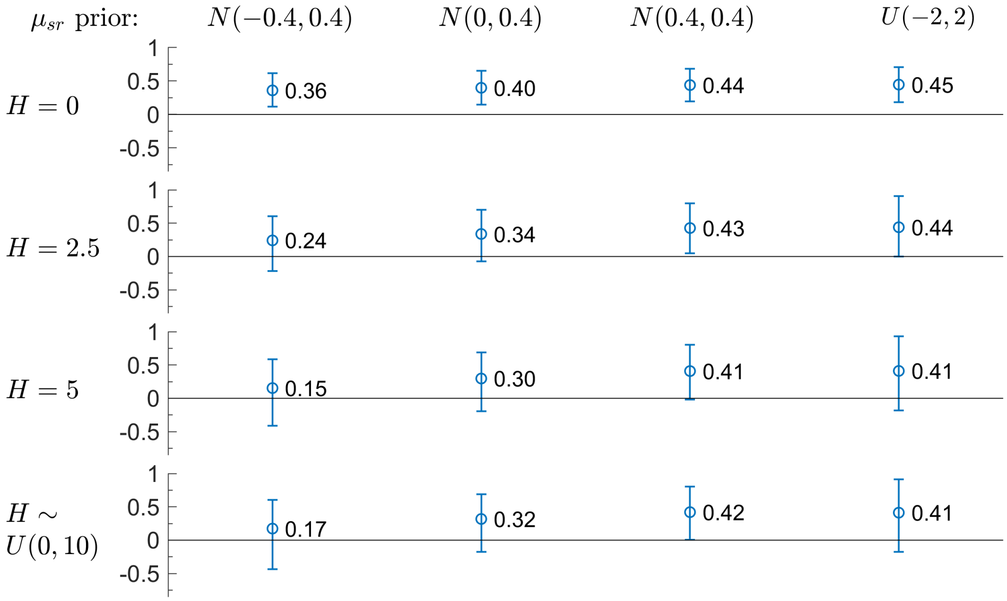

Panel B: Sharpe Ratio for 2020 Conditional on 1963–2019 Returns

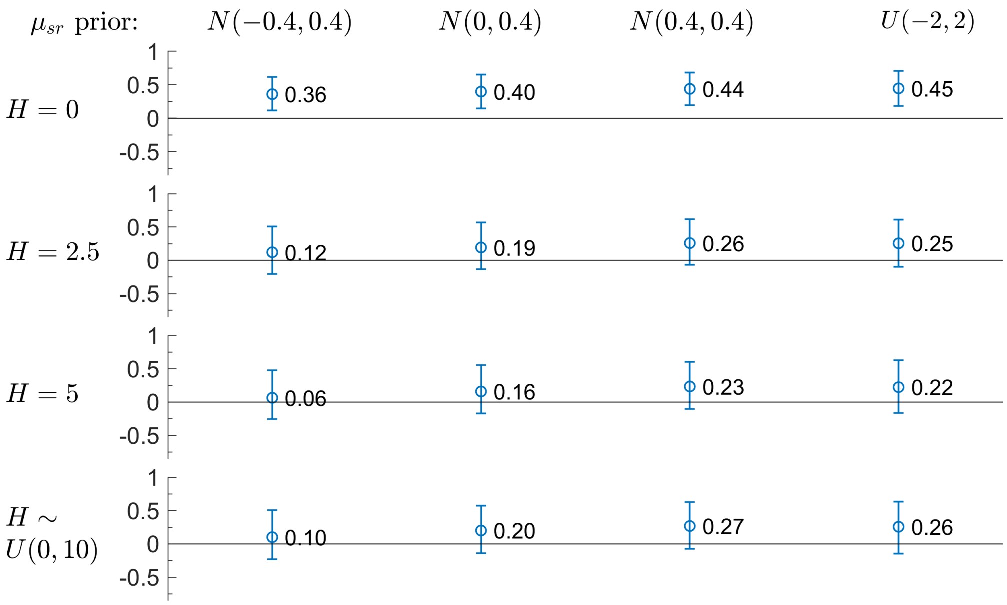

Figure 4: Posterior Sharpe Ratios for Profitability Portfolio

This figure presents posterior beliefs for the Sharpe Ratios of a portfolio formed by sorting stocks on profitability, as defined in the header of Table 1. Panel A presents beliefs about the unconditional Sharpe Ratio, Panels B and C about the conditional Sharpe Ratios in Q1 of 2020 and Q4 of 2029, respectively. Each panel shows posteriors for twelve possible priors, combining three possible priors for the unconditional Sharpe Ratio, μ_{sr}, with four possible priors for the half-life of shocks to conditional Sharpe Ratio, H. All Sharpe Ratios and H are annualized. Our sample consists of 226 quarterly observations from Q3 1963 through Q4 2019.

Panel A: Unconditional Sharpe Ratio

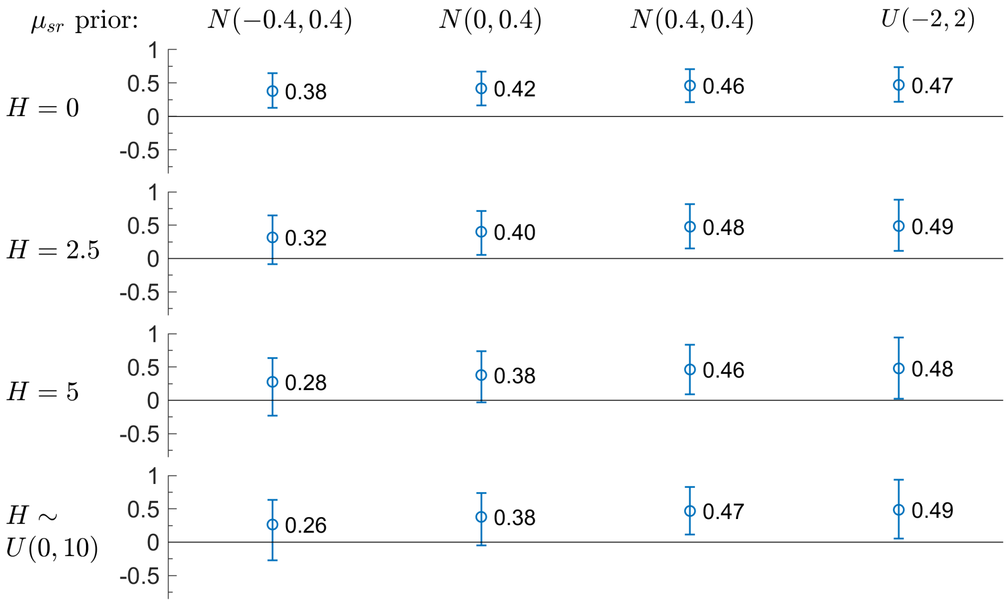

Panel B: Sharpe Ratio for 2020 Conditional on 1963–2019 Returns

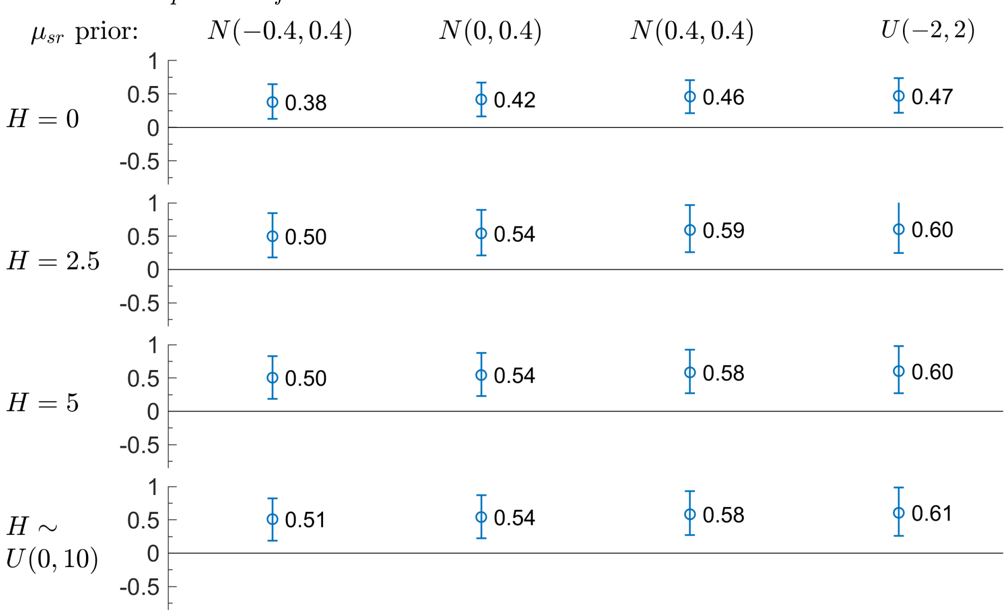

Figure 5: Posterior Sharpe Ratios for Size Portfolio

This figure presents posterior beliefs for the Sharpe Ratios of a portfolio formed by sorting stocks on size, as defined in the header of Table 1. Panel A presents beliefs about the unconditional Sharpe Ratio, Panels B and C about the conditional Sharpe Ratios in Q1 of 2020 and Q4 of 2029, respectively. Each panel shows posteriors for twelve possible priors, combining three possible priors for the unconditional Sharpe Ratio, μ_{sr}, with four possible priors for the half-life of shocks to conditional Sharpe Ratio, H. All Sharpe Ratios and H are annualized. Our sample consists of 226 quarterly observations from Q3 1963 through Q4 2019.

Panel A: Unconditional Sharpe Ratio

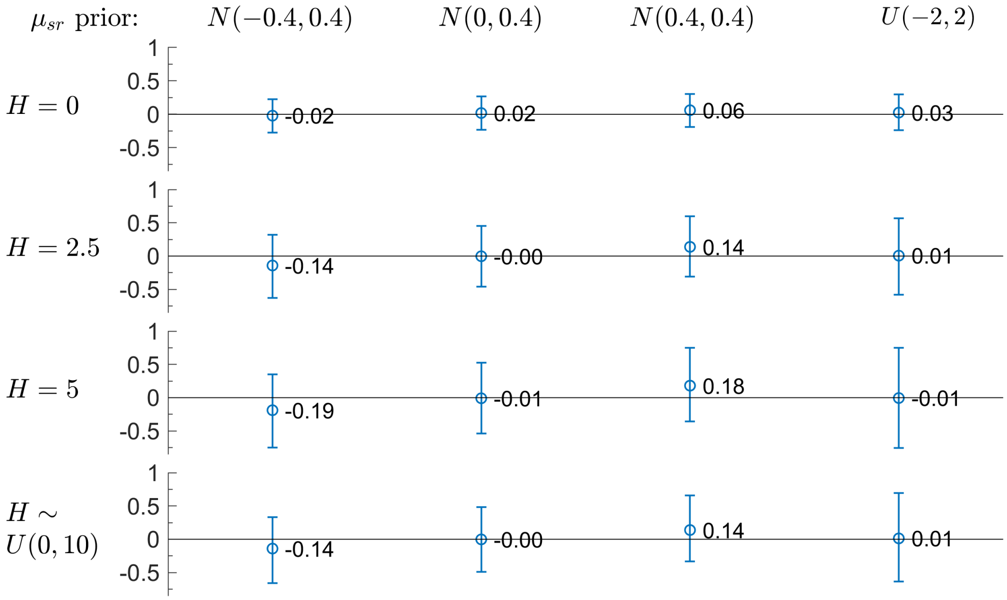

Panel B: Sharpe Ratio for 2020 Conditional on 1963–2019 Returns

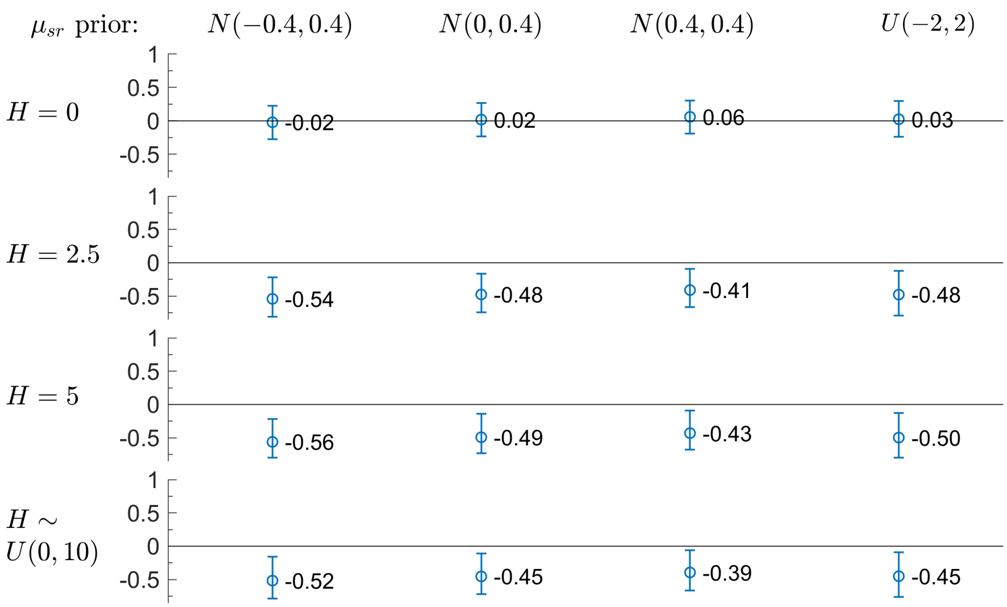

Figure 6: Posteriors on Conditional Sharpe Ratios Across Time

This figure presents posterior beliefs for Sharpe Ratios of value, investment, profitability, and size portfolios, as defined in the header of Table 1. Posterior distributions are formed using the prior that $\mu_{sr} \sim N(0, 0.4)$ and $H \sim U(0, 10)$) and the full sample of data. For each possible parameterization, we compute the expected value of conditional expected returns in each quarter of the sample using the full sample of realized returns. We then plot the average of this quantity across posterior draws for the value portfolio in Panel A, investment in Panel B, profitability in Panel C, and size in Panel D. Our sample consists of 226 quarterly observations from Q3 1963 through Q4 2019.

Panel A: Value

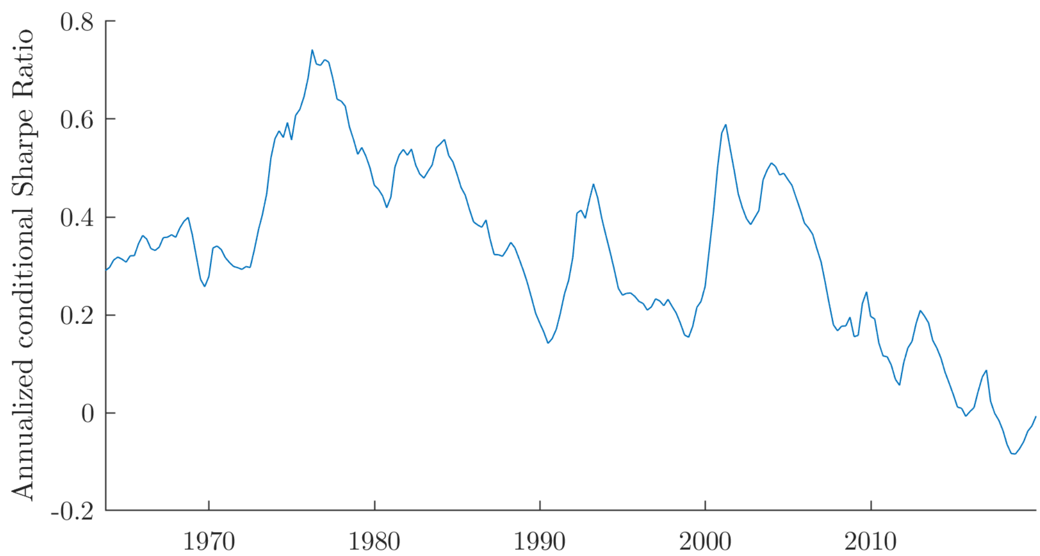

Panel B: Investment

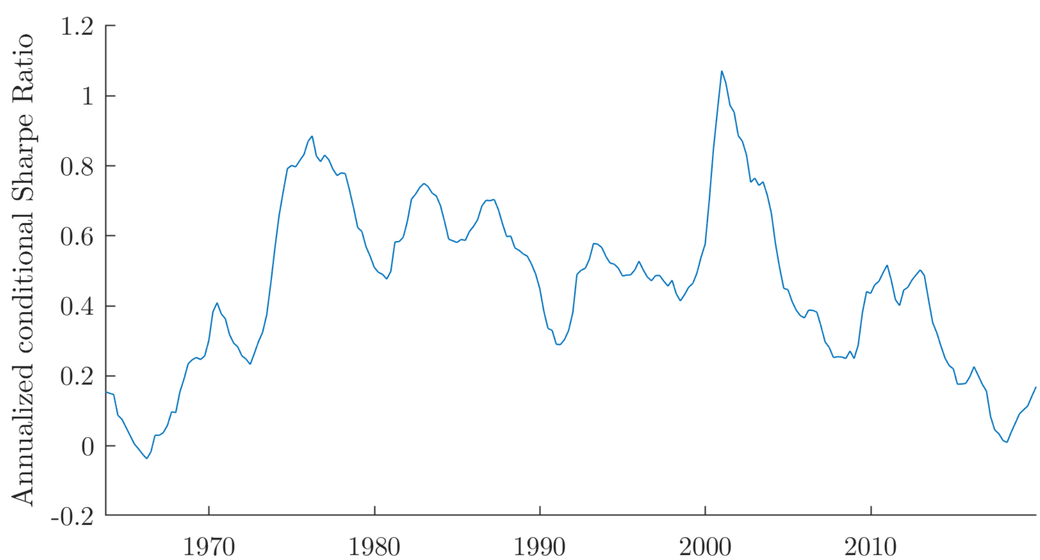

Figure 6: Posteriors on Conditional Sharpe Ratios Across Time (Continued)

Panel C: Profitability

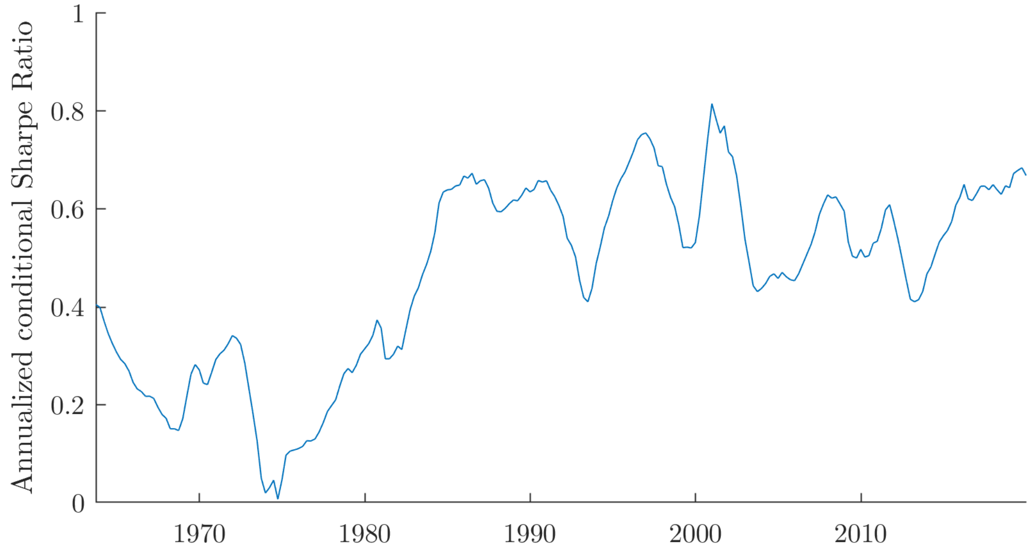

Panel D: Size

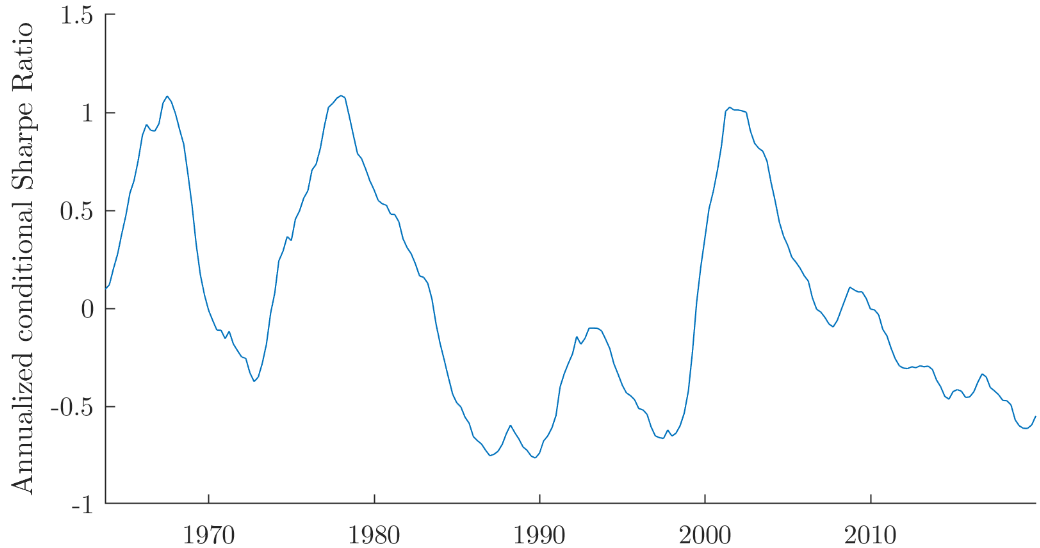

Table 1: Historical Performance of Characteristic-Sorted Portfolios

This table presents statistics summarizing the historical performance of value-weighted quintile portfolios formed on value, investment, profitability, and size characteristics. The value portfolio is based on sorting firms by their book-to-market ratios, the investment portfolio on sorting by the annual growth rate of total assets, the profitability portfolio on sorting by operating profits divided by book equity, and the size portfolio on sorting by market capitalization, as in Fama and French (2015). The investment portfolio is long firms in the lowest quintile and short firms in the highest quintile, while the other three portfolios are long the highest quintile and short the lowest. For each portfolio, we compute market-neutral returns by hedging out market risk using the full-sample market β. We present the mean annualized quarterly return and annualized Sharpe ratio in the full sample, two subsamples, and for the difference between the subsamples. Standard errors based on iid re-sampling of the calendar quarters in our sample are in parenthesis. Our sample consists of 226 quarterly observations from Q3 1963 through Q4 2019.

	Mean (annualized %)				Sharpe Ratio (annualized)			
	All	1963–1991	1992–2019	Diff	All	1964–1990	1991–2019	Diff
Value	4.37	5.52	3.21	-2.30	0.32	0.42	0.22	-0.20
	(1.83)	(2.57)	(2.61)	(3.68)	(0.13)	(0.19)	(0.19)	(0.27)
Investment	4.72	4.25	5.20	0.95	0.44	0.44	0.45	0.02
	(1.41)	(1.99)	(2.01)	(2.84)	(0.13)	(0.19)	(0.19)	(0.27)
Profitability	4.91	3.51	6.33	2.82	0.47	0.36	0.57	0.20
	(1.40)	(1.98)	(1.97)	(2.79)	(0.13)	(0.19)	(0.19)	(0.27)
Size	0.35	1.66	-0.99	-2.66	0.02	0.10	-0.07	-0.17
	(2.02)	(2.82)	(2.87)	(4.02)	(0.14)	(0.19)	(0.19)	(0.27)

Table 2: Autocorrelations of Characteristic-Sorted Portfolios

This table presents statistics summarizing the autocorrelation of value, investment, profitability, and size portfolios, as defined in the header of Table 1. For each characteristic-sorted portfolio, we run overlapping time-series regressions of quarterly returns r_t on a constant and rolling averages of past quarterly returns from months $t-L$ through $t-1$:

$$r_t = a + b_L \left(\frac{1}{L} \sum_{l=1}^{L} r_{t-l} \right) + \epsilon_t, \tag{16}$$

Panel A presents average coefficients b_L in 226-quarter samples simulated under various parameterizations of the model. The parameter γ is the first-order autocorrelation of returns implied by each parameterization, and H is the half-life of shocks to expected returns (in years). Panel B show estimates for the market-neutral portfolios. Standard errors based on iid re-sampling of the calendar months in our sample are in parenthesis. The joint significance row presents the fraction of simulations for which the sum of the $\hat{b}$ across the four portfolios exceeds the sum in observed data. Our sample consists of 226 quarterly observations from Q3 1963 through Q4 2019.

Panel A: Model Simulations

Parameterization			Avg. Coefficients [95% Conf. Interval]		
H	γ	σ(Cond. Sharpe)	b_4	b_{20}	b_{40}
IID		0.00	-0.03	-0.20	-0.49
			[-0.32,0.22]	[-1.08,0.38]	[-2.14,0.45]
2.5	2.5%	0.12	0.04	0.01	-0.21
			[-0.24,0.30]	[-0.79,0.53]	[-1.65,0.61]
2.5	5%	0.22	0.11	0.13	-0.05
			[-0.19,0.36]	[-0.62,0.61]	[-1.38,0.69]
2.5	10%	0.38	0.22	0.30	0.12
			[-0.08,0.48]	[-0.36,0.71]	[-1.02,0.78]
5	2.5%	0.15	0.04	0.02	-0.15
			[-0.25,0.29]	[-0.79,0.56]	[-1.61,0.66]
5	5%	0.26	0.10	0.17	0.03
			[-0.20,0.00]	[-0.61,0.37]	[-1.31,0.66]
10	2.5%	0.12	0.09	0.02	-0.00
			[0.00,0.00]	[-0.27,0.28]	[-0.86,0.56]
10	5%	0.20	0.15	0.07	0.14
			[0.00,0.00]	[-0.23,0.35]	[-0.68,0.67]

Panel B: Estimates for historical data

	b_4	b_{20}	b_{40}
Value	0.19	0.05	0.24
Investment	0.21	-0.29	-0.73
Profitability	0.21	-0.24	0.16
Size	0.47	0.20	0.34
Pooled	0.30	0.06	0.18
iid p-value	0.0%	24.5%	4.7%
Pooled (without size)	0.20	-0.13	-0.01
iid p-value	0.1%	45.1%	15.2%

Table 3: OLS and GLS with Time-Varying Expected Returns

This table presents estimates of unconditional expected returns of value, investment, profitability, and size portfolios, as defined in the header of Table 1, under a variety of assumptions about the magnitude and persistence of variations in conditional expected returns. The model-implied autocorrelation structure of returns are summarized by H, the half-life of shocks to expected returns (in years), and γ, the first-order autocorrelation of realized returns. We estimate unconditional expected returns using both OLS and GLS. We calculate standard errors using Newey and West (1987) and the model-implied correlation matrix of the regression error terms. Our sample consists of 226 quarterly observations from Q3 1963 through Q4 2019.

Panel A: Value

	H (years)	0	2.5			5		10	
	γ (%)	0	2.5	5	10	2.5	5	2.5	5
OLS	$\hat{\mu}$ (%)	4.37	4.37	4.37	4.37	4.37	4.37	4.37	4.37
	Newey-West	(1.83)	(1.86)	(1.86)	(1.86)	(1.86)	(1.86)	(1.81)	(1.81)
	Model	(1.83)	(2.39)	(2.84)	(3.57)	(2.77)	(3.46)	(3.27)	(4.24)
GLS	$\hat{\mu}$ (%)	4.37	4.19	4.07	3.92	3.99	3.78	3.79	3.45
	Model	(1.83)	(2.39)	(2.83)	(3.55)	(2.75)	(3.42)	(3.24)	(4.17)

Panel B: Investment

	H (years)	0	2.5			5		10	
	γ (%)	0	2.5	5	10	2.5	5	2.5	5
OLS	$\hat{\mu}$ (%)	4.72	4.72	4.72	4.72	4.72	4.72	4.72	4.72
	Newey-West	(1.41)	(1.63)	(1.63)	(1.63)	(1.55)	(1.55)	(1.46)	(1.46)
	Model	(1.42)	(1.84)	(2.19)	(2.76)	(2.14)	(2.67)	(2.52)	(3.27)
GLS	$\hat{\mu}$ (%)	4.72	4.52	4.41	4.30	4.28	4.05	4.01	3.62
	Model	(1.42)	(1.84)	(2.18)	(2.74)	(2.13)	(2.64)	(2.50)	(3.22)

Panel C: Profitability

	H (years)	0	2.5			5		10	
	γ (%)	0	2.5	5	10	2.5	5	2.5	5
OLS	$\hat{\mu}$ (%)	4.91	4.91	4.91	4.91	4.91	4.91	4.91	4.91
	Newey-West	(1.39)	(1.47)	(1.47)	(1.47)	(1.32)	(1.32)	(1.41)	(1.41)
	Model	(1.39)	(1.81)	(2.15)	(2.70)	(2.10)	(2.62)	(2.48)	(3.21)
GLS	$\hat{\mu}$ (%)	4.91	4.96	5.02	5.12	4.94	5.02	4.87	4.94
	Model	(1.39)	(1.81)	(2.14)	(2.68)	(2.09)	(2.59)	(2.45)	(3.16)

Panel D: Size

	H (years)	0	2.5			5		10	
	γ (%)	0	2.5	5	10	2.5	5	2.5	5
OLS	$\hat{\mu}$ (%)	0.35	0.35	0.35	0.35	0.35	0.35	0.35	0.35
	Newey-West	(2.00)	(2.96)	(2.96)	(2.96)	(3.15)	(3.15)	(2.83)	(2.83)
	Model	(2.00)	(2.61)	(3.10)	(3.90)	(3.02)	(3.78)	(3.57)	(4.63)
GLS	$\hat{\mu}$ (%)	0.35	0.26	0.17	0.00	0.26	0.15	0.33	0.21
	Model	(2.00)	(2.61)	(3.09)	(3.87)	(3.01)	(3.74)	(3.54)	(4.55)

Table 4: Maximum Likelihood Hypothesis Tests

This table presents parameter estimates and hypothesis tests based on maximum-likelihood estimates of our model for value, investment, profitability, and size portfolios, as defined in the header of Table 1. Panel A presents estimates of μ, the unconditional expected return; σ_r, the unconditional standard deviation of returns; H, the half-life of time-variations in conditional mean returns; and γ, the first-order autocorrelation of factor returns. μ, σ_r, and H are all annualized. The rows labelled $\mu = 0$ re-estimate the model with μ restricted to zero, and give a p-value for this restriction based on a likelihood ratio test. Panel B presents a variety of hypothesis tests for different restrictions on H and γ. For each restriction, we present the estimated γ, and likelihood-ratio p-values for both the H and γ restrictions (Rest. p-value) and $\mu = 0$ given the H and γ restrictions. Our sample consists of 226 quarterly observations from Q3 1963 through Q4 2019.

Panel A: Tests for $\mu = 0$

	IID			Time-varying means				
	μ (%)	σ_r (%)	p-value (%)	μ (%)	σ_r (%)	H (years)	γ (%)	p-value (%)
Value	4.37	13.76		4.60	13.75	14.87	-0.58	
$\mu = 0$	0.00	13.93	1.8%	0.00	13.90	20.00	2.73	10.2%
Investment	4.72	10.62		5.19	10.61	9.38	-0.92	
$\mu = 0$	0.00	10.88	0.1%	0.00	10.82	14.66	5.25	5.7%
Profitability	4.91	10.42		4.96	10.42	20.00	-0.43	
$\mu = 0$	0.00	10.70	0.0%	0.00	10.62	20.00	4.46	5.7%
Size	0.35	15.02		0.31	15.02	1.33	15.44	
$\mu = 0$	0.00	15.03	86.3%	0.00	15.04	2.16	14.84	37.8%

Panel B: Restrictions on H and γ

Restrictions	H (years):		2.5			5			10		
	γ (%):	IID	-	2.5	5	-	2.5	5	-	2.5	5
Value	Rest. p-value (%)	45.2	45.2	60.7	40.1	46.0	51.4	31.2	73.9	46.1	28.4
	$\mu = 0$ p-value (%)	1.8	7.1	8.1	15.4	9.5	14.9	27.6	9.6	24.5	42.0
Investment	Rest. p-value (%)	30.3	41.3	71.5	65.6	30.3	51.6	40.0	95.1	39.1	28.6
	$\mu = 0$ p-value (%)	0.1	6.3	1.5	4.5	8.9	4.5	13.0	5.5	11.2	27.4
Profitability	Rest. p-value (%)	54.8	55.8	68.8	44.8	55.9	61.5	35.5	54.8	58.1	34.5
	$\mu = 0$ p-value (%)	0.0	1.1	0.6	2.1	2.0	1.9	5.6	4.2	4.9	12.3
Size	Rest. p-value (%)	0.0	29.2	2.2	8.3	5.3	1.0	2.6	1.2	0.5	0.8
	$\mu = 0$ p-value (%)	86.3	100.0	92.2	97.5	100.0	93.1	100.0	100.0	92.6	96.8

Table 5: Priors

This table presents summary statistics for the variety of prior distributions we use for our Bayesian estimates. Panel A lists the possible priors we consider for μ_{sr}, the annualized unconditional Sharpe ratio; H, the half life of variations in conditional Sharpe ratios in years; σ_r, annualized unconditional return volatility; σ_{sr}, the annualized volatility of conditional Sharpe ratios; and ρ, the correlation between shocks to realized returns and conditional Sharpe ratios. N indicates a normal distribution, U indicates a uniform distribution, and a number indicates a dogmatic prior that the parameter equals a that value. Panel B presents the means and 95% confidence intervals for prior distributions of μ, the annualized unconditional expected return; γ, the first-order autocorrelation of returns; and σ(Cond. Sharpe), the standard deviation across time of Sharpe ratios conditional on all past realized returns. The conditional Sharpe Ratio distribution is based on the joint return distributions implied by each model parameterization and not observed data.

Panel A: Priors on Transformed Parameters

H	μ_{sr}	σ_r (%)	σ_{sr}	ρ
0	$N(-0.4, 0.4)$	$U(10, 20)$	$U(0, 1)$	$U(-0.5, 0)$
2.5	$N(0, 0.4)$			
5	$N(0.4, 0.4)$			
$U(0, 10)$	$U(-2, 2)$			

Panel B: Moments of Priors

		μ (%)		γ (%)		σ(Cond. Sharpe)	
H	μ_{sr}	mean	95% CI	mean	95% CI	mean	95% CI
0	$N(-0.4, 0.4)$	-5.30	[-15.90,5.16]	-	-	-	-
0	$N(0, 0.4)$	-0.01	[-10.47,10.42]	-	-	-	-
0	$N(0.4, 0.4)$	5.32	[-5.10,16.05]	-	-	-	-
0	$U(-2, 2)$	-0.02	[-25.84,25.70]	-	-	-	-
2.5	$N(-0.4, 0.4)$	-5.34	[-16.07,5.07]	4.81	[-0.58,15.44]	0.19	[0.00,0.52]
2.5	$N(0, 0.4)$	-0.04	[-10.42,10.37]	4.76	[-0.59,15.41]	0.19	[0.00,0.52]
2.5	$N(0.4, 0.4)$	5.29	[-5.10,15.98]	4.78	[-0.58,15.43]	0.19	[0.00,0.52]
2.5	$U(-2, 2)$	-0.05	[-25.88,25.72]	4.76	[-0.58,15.45]	0.19	[0.00,0.52]
5	$N(-0.4, 0.4)$	-5.30	[-15.94,5.11]	5.58	[-0.25,16.49]	0.25	[0.00,0.61]
5	$N(0, 0.4)$	0.01	[-10.47,10.47]	5.60	[-0.26,16.47]	0.25	[0.00,0.61]
5	$N(0.4, 0.4)$	5.23	[-5.18,15.94]	5.61	[-0.26,16.54]	0.25	[0.00,0.61]
5	$U(-2, 2)$	0.00	[-25.75,25.93]	5.56	[-0.25,16.41]	0.25	[0.00,0.61]
$U(0, 10)$	$N(-0.4, 0.4)$	-5.32	[-16.03,4.99]	4.91	[-1.91,16.37]	0.23	[0.00,0.63]
$U(0, 10)$	$N(0, 0.4)$	-0.02	[-10.48,10.38]	4.90	[-1.79,16.37]	0.23	[0.00,0.63]
$U(0, 10)$	$N(0.4, 0.4)$	5.29	[-5.14,15.98]	4.92	[-1.76,16.40]	0.23	[0.00,0.63]
$U(0, 10)$	$U(-2, 2)$	-0.01	[-25.81,25.82]	4.90	[-1.83,16.38]	0.23	[0.00,0.63]

Table 6: Posteriors for Value Portfolio

This table presents summary statistics for the posterior distributions of our model's parameters given data for the value portfolio, as described in the header of Table 1. μ is the unconditional mean return, in annualized percent. σ_r is the unconditional volatility of returns, in annualized percent. H is the half life of variations in conditional Sharpe Ratios, in years. γ is the first-order autocorrelation in realized returns implied by the model's parameters, in percent. The twelve possible priors are summarized in Table 5. We present the mean and 95% confidence intervals for each parameter and prior. We estimate the model using an M.C.M.C. procedure described in the Appendix. Our sample consists of 226 quarterly observations from Q3 1963 through Q4 2019.

Prior		μ		σ_r		H		σ(Cond. Sharpe)	
μ_{sr}	H (years)	mean	95% CI	mean	95% CI	mean	95% CI	mean	95% CI
0	$N(-0.4, 0.4)$	3.41	[-0.10,6.73]	12.69	[12.69,14.99]	-	-	-	-
0	$N(0, 0.4)$	3.94	[0.55,7.37]	12.62	[12.62,14.95]	-	-	-	-
0	$N(0.4, 0.4)$	4.47	[1.04,7.89]	12.58	[12.58,14.98]	-	-	-	-
0	$U(-2, 2)$	4.35	[0.72,8.02]	12.56	[12.56,14.98]	-	-	-	-
2.5	$N(-0.4, 0.4)$	2.59	[-2.95,6.92]	12.74	[12.74,15.30]	-	-	0.14	[0.00,0.46]
2.5	$N(0, 0.4)$	3.45	[-1.40,7.70]	12.70	[12.70,15.25]	-	-	0.13	[0.00,0.45]
2.5	$N(0.4, 0.4)$	4.42	[-0.12,8.84]	12.69	[12.69,15.26]	-	-	0.13	[0.00,0.44]
2.5	$U(-2, 2)$	4.18	[-1.09,9.19]	12.67	[12.67,15.27]	-	-	0.13	[0.00,0.45]
5	$N(-0.4, 0.4)$	1.72	[-4.93,6.81]	12.74	[12.74,15.36]	-	-	0.19	[0.00,0.56]
5	$N(0, 0.4)$	3.16	[-2.65,7.89]	12.72	[12.72,15.32]	-	-	0.16	[0.00,0.54]
5	$N(0.4, 0.4)$	4.33	[-0.75,9.26]	12.69	[12.69,15.32]	-	-	0.15	[0.00,0.52]
5	$U(-2, 2)$	4.03	[-2.68,10.05]	12.71	[12.71,15.33]	-	-	0.17	[0.00,0.54]
$U(0, 10)$	$N(-0.4, 0.4)$	1.86	[-5.33,6.78]	12.78	[12.78,15.38]	5.12	[0.35,9.77]	0.18	[0.00,0.58]
$U(0, 10)$	$N(0, 0.4)$	3.28	[-2.42,7.86]	12.70	[12.70,15.32]	4.84	[0.34,9.75]	0.16	[0.00,0.54]
$U(0, 10)$	$N(0.4, 0.4)$	4.38	[-0.62,9.21]	12.67	[12.67,15.30]	4.78	[0.30,9.73]	0.15	[0.00,0.53]
$U(0, 10)$	$U(-2, 2)$	4.01	[-2.90,9.64]	12.67	[12.67,15.31]	4.80	[0.26,9.74]	0.15	[0.00,0.56]

Table 7: Posteriors for Investment Portfolio

This table presents summary statistics for the posterior distributions of our model's parameters given data for the investment portfolio, as described in the header of Table 1. μ is the unconditional mean return, in annualized percent. σ_r is the unconditional volatility of returns, in annualized percent. H is the half life of variations in conditional Sharpe Ratios, in years. γ is the first-order autocorrelation in realized returns implied by the model's parameters, in percent. The twelve possible priors are summarized in Table 5. We present the mean and 95% confidence intervals for each parameter and prior. We estimate the model using an M.C.M.C. procedure described in the Appendix. Our sample consists of 226 quarterly observations from Q3 1963 through Q4 2019.

Prior		μ		σ_r		H		σ(Cond. Sharpe)	
μ_{sr}	H (years)	mean	95% CI	mean	95% CI	mean	95% CI	mean	95% CI
0	$N(-0.4, 0.4)$	3.83	[1.27,6.56]	9.76	[9.76,11.61]	-	-	-	-
0	$N(0, 0.4)$	4.24	[1.60,6.92]	9.71	[9.71,11.59]	-	-	-	-
0	$N(0.4, 0.4)$	4.66	[2.06,7.26]	9.73	[9.73,11.56]	-	-	-	-
0	$U(-2, 2)$	4.74	[2.01,7.51]	9.75	[9.75,11.55]	-	-	-	-
2.5	$N(-0.4, 0.4)$	2.51	[-2.25,6.26]	9.87	[9.87,11.84]	-	-	0.22	[0.00,0.51]
2.5	$N(0, 0.4)$	3.49	[-0.76,7.15]	9.81	[9.81,11.83]	-	-	0.20	[0.00,0.50]
2.5	$N(0.4, 0.4)$	4.41	[0.47,8.11]	9.80	[9.80,11.81]	-	-	0.19	[0.00,0.50]
2.5	$U(-2, 2)$	4.52	[-0.02,9.27]	9.82	[9.82,11.79]	-	-	0.20	[0.00,0.50]
5	$N(-0.4, 0.4)$	1.62	[-4.25,6.07]	9.90	[9.90,11.91]	-	-	0.27	[0.00,0.60]
5	$N(0, 0.4)$	3.11	[-2.03,7.15]	9.83	[9.83,11.87]	-	-	0.23	[0.00,0.58]
5	$N(0.4, 0.4)$	4.27	[-0.20,8.30]	9.81	[9.81,11.83]	-	-	0.21	[0.00,0.58]
5	$U(-2, 2)$	4.28	[-1.88,9.44]	9.83	[9.83,11.87]	-	-	0.23	[0.00,0.59]
$U(0, 10)$	$N(-0.4, 0.4)$	1.81	[-4.54,6.29]	9.86	[9.86,11.92]	5.15	[0.59,9.76]	0.26	[0.00,0.61]
$U(0, 10)$	$N(0, 0.4)$	3.29	[-1.81,7.11]	9.84	[9.84,11.83]	4.55	[0.36,9.66]	0.22	[0.00,0.59]
$U(0, 10)$	$N(0.4, 0.4)$	4.34	[0.04,8.23]	9.79	[9.79,11.84]	4.49	[0.35,9.66]	0.20	[0.00,0.58]
$U(0, 10)$	$U(-2, 2)$	4.25	[-1.84,9.27]	9.81	[9.81,11.83]	4.55	[0.38,9.76]	0.22	[0.00,0.59]

Table 8: Posteriors for Profitability Portfolio

This table presents summary statistics for the posterior distributions of our model's parameters given data for the profitability portfolio, as described in the header of Table 1. μ is the unconditional mean return, in annualized percent. σ_r is the unconditional volatility of returns, in annualized percent. H is the half life of variations in conditional Sharpe Ratios, in years. γ is the first-order autocorrelation in realized returns implied by the model's parameters, in percent. The twelve possible priors are summarized in Table 5. We present the mean and 95% confidence intervals for each parameter and prior. We estimate the model using an M.C.M.C. procedure described in the Appendix. Our sample consists of 226 quarterly observations from Q3 1963 through Q4 2019.

Prior		μ		σ_r		H		σ(Cond. Sharpe)	
μ_{sr}	H (years)	mean	95% CI	mean	95% CI	mean	95% CI	mean	95% CI
0	$N(-0.4, 0.4)$	3.98	[1.36,6.70]	9.63	[9.63,11.39]	-	-	-	-
0	$N(0, 0.4)$	4.37	[1.72,6.94]	9.57	[9.57,11.38]	-	-	-	-
0	$N(0.4, 0.4)$	4.81	[2.28,7.40]	9.53	[9.53,11.33]	-	-	-	-
0	$U(-2, 2)$	4.92	[2.27,7.59]	9.54	[9.54,11.36]	-	-	-	-
2.5	$N(-0.4, 0.4)$	3.25	[-0.88,6.65]	9.66	[9.66,11.60]	-	-	0.15	[0.00,0.47]
2.5	$N(0, 0.4)$	4.12	[0.53,7.27]	9.61	[9.61,11.56]	-	-	0.13	[0.00,0.45]
2.5	$N(0.4, 0.4)$	4.88	[1.53,8.22]	9.61	[9.61,11.55]	-	-	0.13	[0.00,0.45]
2.5	$U(-2, 2)$	4.98	[1.18,8.90]	9.63	[9.63,11.56]	-	-	0.13	[0.00,0.46]
5	$N(-0.4, 0.4)$	2.87	[-2.28,6.49]	9.71	[9.71,11.63]	-	-	0.18	[0.00,0.55]
5	$N(0, 0.4)$	3.92	[-0.32,7.51]	9.65	[9.65,11.61]	-	-	0.16	[0.00,0.53]
5	$N(0.4, 0.4)$	4.77	[0.92,8.48]	9.61	[9.61,11.60]	-	-	0.15	[0.00,0.52]
5	$U(-2, 2)$	4.93	[0.23,9.65]	9.62	[9.62,11.61]	-	-	0.16	[0.00,0.54]
$U(0, 10)$	$N(-0.4, 0.4)$	2.71	[-2.78,6.52]	9.67	[9.67,11.67]	5.09	[0.38,9.82]	0.19	[0.00,0.58]
$U(0, 10)$	$N(0, 0.4)$	3.89	[-0.51,7.51]	9.62	[9.62,11.64]	4.77	[0.37,9.75]	0.16	[0.00,0.54]
$U(0, 10)$	$N(0.4, 0.4)$	4.79	[1.19,8.42]	9.60	[9.60,11.57]	4.61	[0.25,9.72]	0.14	[0.00,0.51]
$U(0, 10)$	$U(-2, 2)$	4.96	[0.56,9.47]	9.60	[9.60,11.56]	4.66	[0.44,9.70]	0.15	[0.00,0.53]

Table 9: Posteriors for Size Portfolio

This table presents summary statistics for the posterior distributions of our model's parameters given data for the size portfolio, as described in the header of Table 1. μ is the unconditional mean return, in annualized percent. σ_r is the unconditional volatility of returns, in annualized percent. H is the half life of variations in conditional Sharpe Ratios, in years. γ is the first-order autocorrelation in realized returns implied by the model's parameters, in percent. The twelve possible priors are summarized in Table 5. We present the mean and 95% confidence intervals for each parameter and prior. We estimate the model using an M.C.M.C. procedure described in the Appendix. Our sample consists of 226 quarterly observations from Q3 1963 through Q4 2019.

Prior		μ		σ_r		H		σ(Cond. Sharpe)	
μ_{sr}	H (years)	mean	95% CI	mean	95% CI	mean	95% CI	mean	95% CI
0	$N(-0.4, 0.4)$	-1.07	[-4.63,2.46]	13.78	[13.78,16.33]	-	-	-	-
0	$N(0, 0.4)$	0.26	[-3.23,3.76]	13.78	[13.78,16.36]	-	-	-	-
0	$N(0.4, 0.4)$	1.55	[-1.75,4.99]	13.77	[13.77,16.36]	-	-	-	-
0	$U(-2, 2)$	0.38	[-3.59,4.48]	13.77	[13.77,16.33]	-	-	-	-
2.5	$N(-0.4, 0.4)$	-3.15	[-8.57,2.02]	13.81	[13.81,16.55]	-	-	0.42	[0.19,0.56]
2.5	$N(0, 0.4)$	-0.02	[-5.07,5.06]	13.78	[13.78,16.56]	-	-	0.41	[0.17,0.56]
2.5	$N(0.4, 0.4)$	3.21	[-1.91,8.39]	13.78	[13.78,16.59]	-	-	0.41	[0.19,0.57]
2.5	$U(-2, 2)$	0.09	[-7.95,7.92]	13.81	[13.81,16.59]	-	-	0.42	[0.18,0.57]
5	$N(-0.4, 0.4)$	-3.90	[-9.83,1.92]	13.97	[13.97,16.81]	-	-	0.49	[0.19,0.65]
5	$N(0, 0.4)$	-0.05	[-5.76,5.63]	13.96	[13.96,16.78]	-	-	0.47	[0.14,0.64]
5	$N(0.4, 0.4)$	3.83	[-1.94,9.68]	13.95	[13.95,16.76]	-	-	0.48	[0.17,0.64]
5	$U(-2, 2)$	-0.08	[-10.61,10.56]	13.96	[13.96,16.79]	-	-	0.48	[0.18,0.64]
$U(0, 10)$	$N(-0.4, 0.4)$	-3.33	[-9.13,2.05]	13.77	[13.77,16.65]	3.57	[0.80,9.10]	0.42	[0.15,0.65]
$U(0, 10)$	$N(0, 0.4)$	-0.10	[-5.26,5.09]	13.75	[13.75,16.66]	3.25	[0.82,8.94]	0.40	[0.14,0.64]
$U(0, 10)$	$N(0.4, 0.4)$	3.29	[-1.90,9.08]	13.79	[13.79,16.71]	3.51	[0.84,9.11]	0.42	[0.15,0.64]
$U(0, 10)$	$U(-2, 2)$	0.18	[-8.91,9.75]	13.76	[13.76,16.66]	3.59	[0.84,9.06]	0.42	[0.15,0.65]