Time-Series Efficient Factors*

Sina Ehsani† Juhani Linnainmaa‡

December 2020

Abstract

We show that factors in prominent asset pricing models are unconditionally minimum-variance inefficient. Factor returns are positively autocorrelated while risk, conditional on past returns, is constant. We derive a transformation that turns an autocorrelated factor into a "time-series efficient" factor. Time-series efficient factors earn 40% higher Sharpe ratios than the original factors and contain all the information found in the original factors. Momentum strategies profit from the inefficiencies in factor returns. An asset pricing model with time-series efficient factors, such as an efficient Fama-French five-factor model, prices momentum. Our results impose novel restrictions on theoretical models of asset returns.

JEL classification: G11, G12, G40

Keywords: Factors; asset pricing models; anomalies; momentum

*We thank Wayne Ferson, Christian Goulding, Serhiy Kozak, and Pete Kyle for helpful comments and suggestions. We are also grateful for the feedback by the seminar participants at University of Maryland, IDC Arison School of Business, Louisiana State University, and Penn State University.

†School of Business, Northern Illinois University. Email: sehsani@niu.edu.

‡Dartmouth College and NBER. Linnainmaa is a consultant to Research Affiliates and Citadel, and a partner at Research Affiliates. Neither Citadel nor Research Affiliates provided any funding for this research. Email: Juhani.T.Linnainmaa@tuck.dartmouth.edu. Mailing address: 100 Tuck Hall, Hanover, NH 03755, United States. Tel: +1 (603) 646-3160.

1 Introduction

The absence of arbitrage guarantees the existence of a positive stochastic discount factor (SDF), m, that satisfies

$$\mathbb{E}_t[m_{t+1}r_{t+1}] = 0 \tag{1}$$

for the excess return r of any asset. The stochastic discount factor representation is equivalent to the linear expected return-beta model if the SDF is linear in the same factors (Dybvig and Ingersoll Jr, 1982; Hansen and Jagannathan, 1991). Multifactor models assume that the stochastic discount factor can be represented by a linear combination of K common factors f_k, which translates into a linear factor model of excess returns:

$$r_{i,t} = \alpha_i + \sum_{k=1}^{K} \beta_{i,k} f_{k,t} + \epsilon_{i,t}. \tag{2}$$

Huberman et al. (1987) and Grinblatt and Titman (1987) show that the linear factor pricing model in equation (2) implies that a combination of factors is minimum-variance efficient (MVE). Therefore, the intercept in *unconditional* time-series regression of excess returns of a portfolio on the k factors is a test of the MVE of the model and the SDF representation. If α_i is zero, an SDF prices the assets and equation (1) is satisfied.

In absence of an economic model, empiricists create proxies for the factors f_k in equation (2) by constructing factor portfolios. The MVE assumption for the factor portfolios is implicit in this practice. When the factors are not *unconditionally* MVE, the asset pricing model may not explain the mean return of asset i using unconditional coefficients. The nonzero alpha may arise in the unconditional regression because of the correlation between asset i's return and the model inefficiencies. We either need to estimate conditional regressions—which is equivalent to increasing the number of factors—or we need unconditionally MVE factors to conduct a valid test of the model by means of time-series regressions. In this paper we factors' time-series properties to form an unconditionally MVE empirical asset pricing model.

We show that all factors in popular asset pricing models are unconditionally inefficient: expected returns are persistent and this predictability is independent of expected variance. If a factor such as value has done well, it typically continues to do well.[1] Positive autocorrelation by itself does not imply that investors can benefit by timing factors. If factors are atypically volatile whenever expected returns are high, a factor-timing strategy may earn Sharpe ratios no different from those of a buy-and-hold strategy. However, if changes in expected returns are not accompanied by proportional changes in expected volatility, factor timing increases Sharpe ratios. We show that while the typical factor's expected return depends on its past return, its variance does not; more of the factor's risk is therefore "unpriced" following periods of low returns. The typical factor is time-series *inefficient*.

We consider the portfolio choice problem of a mean-variance investor who believes that a factor's return follows an AR(1) process to form unconditionally MVE factors. This investor needs just three parameters—the factor's unconditional mean, variance, and autocorrelation—and one variable—the factor's prior return—to compute the optimal allocation that transforms a factor to its time-series *efficient* version. A time-series efficient factor exploits the autocorrelation in factor returns; its weight on the original factor varies to minimize variance while maintaining the expected return. Time-series efficient factors earn higher Sharpe ratios than the original factors and contain all the information found in the original factors.

An investor can estimate, ex ante, the scope and statistical significance of potential efficiency gains: given beliefs about the factor's unconditional moments, the investor can compute a test statistic for the expected improvement in its Sharpe ratio. Consider, for example, the value factor,

[1] McLean and Pontiff (2016), Arnott et al. (2019), and Ehsani and Linnainmaa (2019) show that a large number of factors are positively autocorrelated.

HML. This factor has an annualized mean of 3.9% with a standard deviation of 9.7%. An investor who holds this factor with a constant weight therefore earns a Sharpe ratio of 0.40. However, because this factor is significantly autocorrelated—the first-order autocorrelation is 0.17 (t-value = 4.46)—while HML's conditional risk is constant, an investor can vary the investment on HML to preserve its premium while taking less risk. Using this information alone, we can predict that the factor's unconditional Sharpe ratio should improve by 0.32 units if we were to exhaust HML's first-order autocorrelation. We also predict this improvement to be statistically significant with an expected z-value of 2.54.

How do these predictions compare with the *realized* improvements? HML's actual Sharpe ratio increases by 0.26 units, and this improvement is significant with a z-value of 3.13. In a five-factor model regression, this efficient factor's alpha is 2.0% (t-value = 4.04). By contrast, in an *efficient* five-factor model regression, the original HML's alpha is 0.0% (t-value = 0.04). That is, the efficient HML contains all the information in the standard HML and more. The time-series efficient HML no longer carries an autocorrelation that can be harvested for alpha because our procedure transforms its autocorrelation into a higher Sharpe ratio.

An investor does not need to know the factors' precise unconditional moments to gain significant efficiency. We show that even large estimation errors in means, variances, and autocorrelations reduce efficient factors' Sharpe ratios relatively little. An investor whose priors imply skepticism about these parameters, but trades under the assumption that factors are at least weakly autocorrelated, gains significantly. Moreover, we show that these moments are quite stable over time for the average factor; investors can learn enough about the required inputs from just a small sliver of data. Consistent with such stability, we find that *real-time* implementations of efficient factors perform just as well as (infeasible) factors that use the original factors' full-sample moments.

Our results are not confined to the Fama-French five-factor model. We examine other factor models—including the four-factor models of Novy-Marx (2013), Hou et al. (2015), and Stambaugh and Yuan (2017)—and find the same pattern: time-series efficient factors dominate standard factors. Time-series efficiency is not the only method for improving mean-variance efficiency. Cohen and Polk (1998), Asness et al. (2000), and Novy-Marx (2013) show that industry-neutral factors typically outperform the standard factors, and Daniel et al. (2020) show that loadings-hedged factors share this property. We show that our methods and these other methods complement each other; an investor can improve already improved factors by conditioning out the autocorrelations. The squared Sharpe ratio of the tangency portfolio of factors in Daniel et al. (2020), for example, is 2.24, compared to 1.15 of the original Fama-French five-factor model. The squared Sharpe ratio of the optimal combination of time-series efficient versions of Daniel et al.'s (2020) factors is 2.45.

Time-series efficiency relates to momentum in individual stock returns (Jegadeesh and Titman, 1993). Fama (2014), for example, argues that "Most prominent [anomaly] is the momentum in short-term returns documented by Jegadeesh and Titman (1993), which is a problem for all asset pricing models that do not add exposure to momentum as an explanatory factor, and which in my view is the biggest challenge to market efficiency." We show that adding momentum is not necessary. The momentum factor serves a purpose only when the original factors are time-series *inefficient*. Instead of adding the momentum factor, we can replace the Fama-French factors with their efficient counterparts and capture momentum without adding a distinct momentum factor.

A comparison between the standard and time-series efficient five-factor models illustrates the connection between factor unconditional efficiency and momentum. Carhart's (1997) UMD factor earns a monthly CAPM alpha of 73 basis points (t-value = 4.48) from 1963 through 2018, and its five-factor model alpha is 72 basis points (t-value = 4.44).[2] UMD's alpha in the *efficient* five-factor model, by contrast, is just 16 basis points (t-value = 1.04). The reason for this improvement is that

[2] Fama and French (2016) also find that "[n]o combination of the factors in [the five-factor model] explains average returns on portfolios formed on momentum."

the predictable time-series variations in the factor premiums in the Fama-French model alone are responsible for a significant share of UMD's profits; therefore, when we exhaust the momentum in factor returns, UMD turns largely redundant. Time-series efficient factor models push momentum back into the factors.

Our results are important from the viewpoint of investing. Because the momentum factor's five-factor model alpha is statistically significant, an investor who trades that model's five factors—the market, size, value, profitability, and investment factors—can earn a higher Sharpe ratio by trading also momentum. The investor can, however, "trade momentum" in three distinct ways. The first method is to trade momentum in the cross section of stocks returns (Jegadeesh and Titman, 1993). The second method is to trade "factor momentum" in addition to the five factors: invest in the factors that have done well and shun those that have done poorly (Ehsani and Linnainmaa, 2019). The third method, based on our results, is that investors can distribute factor momentum back into the factors; by trading efficient factors, the investor need not be concerned with momentum.

Why might efficient factors have these properties? Suppose that we have an unobserved factor F_t for which we attempt to create a mimicking portfolio. The mimicking portfolio $\hat{F}_t$, such as a HML portfolio sorted on characteristics, will contain an error: $\hat{F}_t = F_t + \hat{\varepsilon}_t$, where $\hat{\varepsilon}_t$ is the noise unrelated to F_t. That is, there is no reason why, for example, a portfolio sort with fixed 30th and 70th percentile breakpoints would be the least contaminated proxy for F_t. The actual factor F_t might appear to relate weakly to our factor $\hat{F}_t$, but only because of the measurement error. By constructing time-series efficient factors, we minimize $\hat{\varepsilon}_t$ and get closer to the true factors.

Our results are informative about factors' risk-return dynamics. Throughout this study, we use simulated data side by side with the actual data. In these simulations *latent* factor returns are autocorrelated and investors can trade noisy versions of these latent factors—an assumption that parallels the idea that HML is a noisy proxy for the true value factor. We simulate data from three economies with different assumptions about how factor risk responds to past returns: it either remains constant, or increases or decreases in past returns. We show that the results from all our tests come closest to the economy where risk and return are independent from one another. If factors were riskier when their expected returns are high, the gains from time-series efficiency would be lower than what they actually are; and if risks were to fall when expected returns rise, the gains would be far larger than what they actually are.

2 Factor returns and time-series predictability

Table 1 shows the average returns and standard deviations of the factors in the Fama-French five-factor model, and reports on the predictability of these factors' means and variances. These factors' average returns range from 2.9 (size) to 6.2 percent per year (market), and the t-values from 2.04 (size) to 3.67 (investment). Much of our analysis is about mean-variance efficiency, which relates to factors' Sharpe ratios. The factors' Sharpe ratios, which are proportional to the t-values, range from 0.27 (size) to 0.49 (investment).

The estimates in Table 1 mask the fact that factor returns are significantly predictable. Ehsani and Linnainmaa (2019), for example, show that most factors are positively autocorrelated: the average factor's return is typically significantly higher, both statistically and economically, after it has performed well.

The left columns of Panel B illustrate this time-series predictability by assigning factors into two groups based on their month $t-1$ returns. The estimates show, for example, that when the size factor's return in month $t-1$ is low, its average annualized return in month t is -1.8%. If its prior-month return is high, its average return in month t is 7.5%. The 9.3% difference in average returns is significant with a t-value of 3.33. Differences for four of the five factors are statistically significant at the 1% level; although the up-minus-down difference for the market is 4.3%, this

Table 1: Fama-French five-factor model: Average returns, standard deviations, and predictability

Panel A reports annualized means, standard deviations, Sharpe ratios, and t-values associated with the average returns for the five factors of the Fama and French (2015) model. In Panel B we assign factors into two groups based on their month $t-1$ returns and report either average annualized returns or volatilities for each group. We use the full sample to create two groups of equal size. We compute the estimates for the "Average" factor in two steps. In the first step we compute month-t average returns separately for the factors assigned into the high and low groups based on their month $t-1$ returns. In the second step we compute time-series averages for the low and high groups and their difference. The high or low group in month t can be empty if no factor gets assigned into it. There are five (nine) months in which all factors are in the low (high) group. We report the high-minus-low difference in average returns and the high/low ratio of variances. We report t-values in parentheses and, for the variance ratio, the p-value associated with the F-test in square brackets. The data for monthly factor returns begin in July 1963 and end in December 2018.

Panel A: Factor means and standard deviations

	Factor				
	MKTRF	SMB	HML	RMW	CMA
Mean	6.17	2.87	3.91	2.98	3.41
Standard deviation	15.22	10.48	9.70	7.47	6.92
Sharpe ratio	0.41	0.27	0.40	0.40	0.49
t-value	3.02	2.04	3.01	2.97	3.67

Panel B: Returns and volatilities conditional on month $t-1$ returns

	Predicting returns			Predicting volatility		
	Prior-month return			Prior-month return		H/L variance
Factor	Low	High	H−L	Low	High	ratio
Market	4.02	8.33	4.31	17.73	12.18	0.47
	(1.19)	(3.60)	(1.05)			[0.00]
Size	−1.77	7.54	9.31	10.38	10.43	1.01
	(−0.90)	(3.80)	(3.33)			[0.93]
Value	−0.79	8.64	9.43	9.37	9.85	1.10
	(−0.45)	(4.61)	(3.65)			[0.36]
Profitability	−1.05	7.10	8.15	7.66	7.09	0.86
	(−0.73)	(5.24)	(4.11)			[0.16]
Investment	−0.09	6.92	7.02	6.40	7.28	1.29
	(−0.08)	(5.00)	(3.81)			[0.02]
Average	−0.25	7.47	7.72	8.03	6.90	0.74
	(−0.23)	(8.04)	(5.36)			[0.00]

difference has a t-value of just 1.05. The last row shows that the average factor's annualized return is -0.3% when it is in the low regime and 7.5% when it is in the high regime. The difference is significant with a t-value of 5.36.

Table 1 shows that expected factor returns are significantly autocorrelated. But should a mean-variance investor vary factor allocations with changes in expected returns? The answer depends on the joint dynamics of variances and returns. Return persistence affects portfolio choice only if the changes in expected returns are not offset by proportionate changes in variance.[3]

The right-hand side columns of Panel B of Table 1 report on the relationship between past returns and future variances. The estimates on the first row show that when the market factor's return in month $t-1$ is low, its annualized volatility in month t is 17.7%; following a high-return month, its volatility is 12.2%. The variance ratio of $\left(\frac{0.122}{0.177}\right)^2 = 0.47$ is statistically significantly different from one with a p-value of 0.00. Market is significantly less, rather than more, volatile following high returns. These estimates together with those in the first three columns show that low market returns signal (1) low market returns at (2) twice the risk. For the size, value, and profitability factors, conditional on past returns, variance is approximately constant. These factors' variance ratios are close to one and most are not statistically significantly different from one. Investment is the only factor whose risk and return move in the same direction. The average factor's variance ratio is 0.74. The typical factor is therefore not more volatile when its prior return is high (and its expected return is high).

Table A1 in the appendix shows that the pattern in average return and volatility are similar to the results in Table 1 when we condition on prior one-year returns. The difference in average returns between the high and low states is 4.9% (t-value = 3.34) for the average factor and the high/low variance ratio is 0.95 (p-value = 0.51). Risks therefore remain approximately constant

[3] In Appendix A, we use a two-period investment problem to demonstrate that l allocation over time does not necessarily increase in expected returns.

even when expected returns move.

3 Efficient factors with conditioning information

3.1 Theoretical framework

A mean-variance investor uses estimates of asset returns to minimize variance for a given expected return. What are the weights of the optimal portfolio when an investor has conditioning information? Ferson and Siegel (2001) study the problem of a Markowitz (1959) investor who uses a return forecasting signal to find the unconditionally minimum variance efficient portfolio.

We use Ferson and Siegel's (2001) framework to construct time-series efficient factors. We first describe this method under the case of one risky asset and the risk-free asset. We then apply this framework to factors with past returns as the conditioning information. We start from a single risky asset with a return of

$$\tilde{R} = \mu(\tilde{S}) + \tilde{\epsilon}, \tag{3}$$

in which $\tilde{R}$ is the risky asset's return in excess of the risk-free rate, $\tilde{S}$ is the predictor (signal), $\mu(\tilde{S})$ is the expected excess return *conditional* on the signal, and $\tilde{\epsilon}$ is the random noise net of the signal with a mean of zero and a variance of $\sigma_\epsilon^2(\tilde{S})$.

The efficient strategy invests $x(\tilde{S})$ in the risky asset and the remainder, $1 - x(\tilde{S})$, in the risk-free asset. This strategy's unconditional expected excess return and variance are

$$\mu_p = \mathbb{E}\big[x(\tilde{S})\mu(\tilde{S})\big], \tag{4}$$

$$\sigma_p^2 = \mathbb{E}\Big[x^2(\tilde{S})\big(\mu^2(\tilde{S}) + \sigma_\epsilon^2(\tilde{S})\big)\Big] - \mu_p^2. \tag{5}$$

Ferson and Siegel (2001) show that the portfolio that minimizes σ_p^2 for a given conditional expectation μ_p invests $x(\tilde{S})$ in the risky asset,

$$x(\tilde{S}) = \frac{\mu_p}{\zeta} \frac{\mu(\tilde{S})}{\mu^2(\tilde{S}) + \sigma_\epsilon^2(\tilde{S})}, \tag{6}$$

in which ζ is a constant,

$$\zeta = \mathbb{E}\left[\frac{\mu^2(\tilde{S})}{\mu^2(\tilde{S}) + \sigma_\epsilon^2(\tilde{S})}\right]. \tag{7}$$

This weighting program produces a unique mean-variance efficient portfolio; no other portfolio has the same unconditional return at a lower unconditional variance (Ferson and Siegel, 2001).

3.2 Efficient factors conditional on past returns

How should investors time factors when the signal $(\tilde{S})$ is a function of past returns? Motivated by the estimates in Table 1, we assume that past returns are related to future returns but not related to variance. Specifically, we assume that returns follows a homeskedastic autoregressive process,

$$\tilde{R}_t = \mu + \rho \tilde{R}_{t-1} + \epsilon_t, \tag{8}$$

$$\text{var}[\epsilon_t | R_{t-1}] = \sigma_\epsilon^2. \tag{9}$$

The factor's conditional expected return under this model is $\mu(\tilde{S}) = \mu + \rho \tilde{R}_{t-1}$. Using equations (4) and (6), the investor's optimal weight on the factor is[4]

$$x(S_t) = \mu_p \frac{SR^2 + 1}{SR^2 + \rho^2} \frac{\mu_p\,(1-\rho) + \rho\, r_{t-1}}{\left(\mu_p\,(1-\rho) + \rho\, r_{t-1}\right)^2 + \sigma_\epsilon^2}, \tag{10}$$

In this formulation μ_p is the factor's unconditional mean, SR is the unconditional Sharpe ratio,

[4]We derive ζ in Appendix B.

ρ is the autocorrelation coefficient, and $\sigma_\epsilon^2 = (1-\rho^2)\sigma^2$ is the constant variance of the noise term. We define *time-series efficient factor* as the portfolio that places the weight $x(S_t)$ from equation (10) on the original factor. A time-series efficient HML, for example, would be the return on a portfolio that optimally times HML given, in this derivation, its month $t-1$ return. In our empirical analyses we use both month $t-1$ return and the average return from month $t-12$ to $t-1$ as the conditioning information.

3.3 An example of the optimal weight function

How much should an investor allocate to a factor when its return has been close to zero or very high or low? To illustrate the optimal policy, we continue with the case in which the signal is the factor's prior-month return. The optimal weight at time t is a function of the factor's mean, standard deviation, and first-order autocorrelation. The average estimates of these monthly parameters and Sharpe ratio for the factors of the Fama-French model from July 1963 through December 2018 are:

Monthly parameter	Mean
Mean	0.32%
Standard deviation	2.85%
Sharpe ratio	0.11
First-order autocorrelation	0.11

The weight function for the "average" factor from equation (10) is then,

$$x(S_t) = x(r_{t-1}) = 13.38 \times \frac{0.32(1-0.11)+0.11r_{t-1}}{\big(0.32(1-0.11)+0.11r_{t-1}\big)^2 + 2.85^2(1-0.11^2)}. \quad (11)$$

Figure 1 shows the optimal weight $x(S_t) = x(r_{t-1})$ as a function of the factor's past return. Because the factor's unconditional premium is positive, an investor does not switch between positive and negative weights at zero; the investor finds it prudent to short the factor only if it has lost more

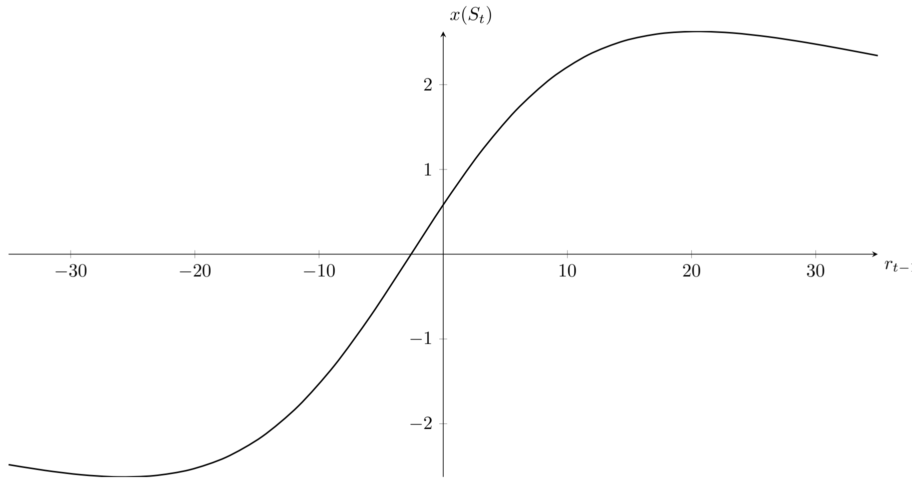

Figure 1: **The optimal weight function.** An investor constructs a time-series efficient factor by predicting month t factor returns with month $t-1$ returns. The optimal weight depends on the factor's mean, standard deviation, and first-order autocorrelation. We compute the average values of these parameters using monthly returns on the five factors of the Fama-French model. This figure plots the optimal weight invested in the factor, $x(S_t) = x(r_{t-1})$, as a function of factors' month $t-1$ return. The returns on the x-axis are in percentage points; the weights on the y-axis in decimals, that is, a value of 1.0 indicates a weight of 100%.

than 2.6% over the prior month.[5]

The weighting program is not monotone in the past return: although the optimal weight initially increases in past return, the optimal strategy begins to scale back as the past return becomes very high. Similarly, the investor begins to scale back on shorting the factor when the factor's past return is very low. The investor's objective of minimizing variance for a given average return drives this nonmonotonicity. If the value of the signal is very high, the investor *could* make an aggressive bet to earn a high expected return; but because the objective is, in effect, to smooth returns, the investor can afford to invest less in the risky asset to lower the strategy's riskiness.[6]

[5]The function is approximately linear for small returns. For example, the slope of the weight is $\left.\frac{dx}{dr}\right|_{r=0\%} = 0.22$ when the factor earns zero and $\left.\frac{dx}{dr}\right|_{r=-2.6\%} = 0.25$ when the optimal weight is zero. The optimal strategy therefore increases allocation to the factor by between 20 and 25% for every unit increase in its last month's return.

[6]Ferson and Siegel (2001, p. 973) describe this intuition as follows:

> "[O]ne could either attain a higher unconditional expected return for the same standard deviation, or, alternatively, attain a lower standard deviation for the same expected return. [. . .] [I]f the objective is to get the smallest unconditional variance for a given average return, then a signal that the expected return is unusually high presents an opportunity to reduce risk by purchasing a smaller amount of the risky asset, while maintaining the portfolio average return."

The optimal factor portfolio earns the target expected return of μ_p with a minimized variance of

$$\sigma_p^2 = \mu_p^2 \left(\frac{1}{\zeta} - 1 \right). \tag{12}$$

We can use this formula to compute the improvement in the Sharpe ratio that an investor can *expect* to attain by switching from the "standard factor" (that is, holding the weight constant at $x = 1$) to the time-series efficient factor. If we create a time-series efficient factor that has the same unconditional mean as the original factor, the improvement in the factor's Sharpe ratio results purely from the lower variance. The expected improvement in Sharpe ratio from moving from the original factor (SR) to the efficient factor (SR^*) is[7]

$$\frac{SR^{*2}}{SR^2} = \frac{1 + \left(\frac{\rho}{SR} \right)^2}{1 - \rho^2}. \tag{13}$$

This equation shows that the expected improvement in the factor's Sharpe ratio depends on the ratio of its autocorrelation coefficient to its Sharpe ratio.[8] When ρ is small, returns are unpredictable, and the ratio in equation (13) converges to one. Because time-series efficiency is about exhausting autocorrelations, the gains evaporate when the time-series predictability vanishes. The potential improvements are the largest for factors that are more autocorrelated, that have lower Sharpe ratios, or that display both of these properties at the same time. The average Fama-French factor earns a Sharpe ratio of 0.11 and has an autocorrelation coefficient of 0.11 in our monthly

They draw a parallel between the optimal policy function and robust estimation; how an agent should make decisions when he or she does not fully trust the model (Hansen and Sargent, 2008). Bekaert and Liu (2004) confirm this conjecture and show that the Ferson and Siegel (2001) solution is equivalent to guarding against misspecifications in the conditional moments.

[7]We derive equation (13) in Appendix B under the assumptions that (1) returns follow an AR(1) process, (2) the first-order autocorrelation is not too large, and (3) returns are homoskedastic.

[8]Campbell and Thompson (2008) derive a similar identity for the proportional increase in expected returns for the quadratic mean-variance investor. The Campbell and Thompson (2008) investor uses the predictive signal to obtain an increase in portfolio return while the Ferson and Siegel (2001) investor uses the predictive signal to minimize portfolio variance.

data. We thus *expect* the the average factor's Sharpe ratio to increase by $\sqrt{\frac{2}{0.99}} = 42\%$.

Equation (13) suggests that the improvements in Sharpe ratios can depend on the frequency of rebalancing. Suppose, for example, that factor autocorrelations at monthly and daily frequencies are similar; then, because daily Sharpe ratio is $1/\sqrt{21}$ of the monthly Sharpe ratio, daily Sharpe ratios would improve more. That is, before transaction costs, high-frequency signals can yield greater improvements in factor efficiency than low-frequency signals; whether or not this happens in the data depends on the behavior of the autocorrelations when we move across frequencies. In our empirical tests we construct time-series efficient factors at both monthly and daily frequencies.

Furthermore, in addition to the expected improvement in a factor's Sharpe ratio, we can also compute the expected z-statistic for the difference between the efficient and original factors' Sharpe ratios. Following Jobson and Korkie (1981), and with the correction from Memmel (2003), the test statistic for the expected difference in Sharpe ratios is

$$z\text{-statistic} = \frac{\sigma_o \mu_e - \sigma_e \mu_o}{\sqrt{\theta}}, \tag{14}$$

$$\text{where } \theta = \frac{1}{T}\Big(2\sigma_e^2\sigma_o^2 - 2\sigma_e\sigma_o\sigma_{e,o} + \frac{1}{2}\mu_e^2\sigma_o^2 + \frac{1}{2}\mu_o^2\sigma_e^2 - \frac{\mu_e\mu_o}{\sigma_e\sigma_o}\sigma_{e,o}^2\Big), \tag{15}$$

where σ_e and σ_o are the standard deviations of the efficient and original factors, μs are the mean returns, and $\sigma_{e,o}$ is the covariance between the efficient and original factors. We derive an approximation for the covariance $\sigma_{e,o}$ in Appendix C.

The predicted z-statistic is a useful tool. Consider, for example, a weakly autocorrelated factor that earns a large Sharpe ratio. Upon implementing the program, we may fail to reject the null hypothesis of no improvement, but only because the large amount of noise in factor returns masks the economically small improvement. The predicted z-statistic takes this lack of power into account. We could solve equations (14) and (15) for the sample size T we expect to require to reach a desired

z-statistic threshold. This computation would tell us whether we have any hope of rejecting the null hypothesis of no improvement under the alternative hypothesis that the program works.

4 Empirical efficient factors

4.1 Time-series efficient Fama-French five-factor model

We start with the Fama and French (2015) five-factor model at monthly frequency to illustrate the properties of time-series efficient factors. After presenting the main results, we construct time-series efficient factors for other popular asset pricing models and at daily frequency.

A time-series efficient factor is a dynamic version of the standard factor; the weight it places on the original factor depends on the factor's prior performance. Because factors are positively autocorrelated, the efficient factor places a higher weight on the standard factor when the factor's own prior return is high; and it shorts the factor when the factor has lost so much money that its conditional risk premium turns negative.

Figure 2 illustrates how the efficient HML changes its weight on the standard HML between August 1963 and December 2018. The efficient factor targets to earn the same expected return as the original factor. We compute the optimal weight from equation (10) using the full-sample estimates of the standard HML's average return, volatility, and first-order serial correlation. In the subsequent tests we compute sample moments recursively to avoid introducing a lookahead bias when comparing the efficient and standard factors.

The optimal weight, which in this analysis derives from HML's month $t-1$ return, has a sample mean of 0.30 but it varies considerably from month to month. The blue line is the efficient factor's average weight on HML over a one-year window around month t. In the late 1990s, for example, when HML underperformed for a prolonged period of time, the optimal factor's weight on HML

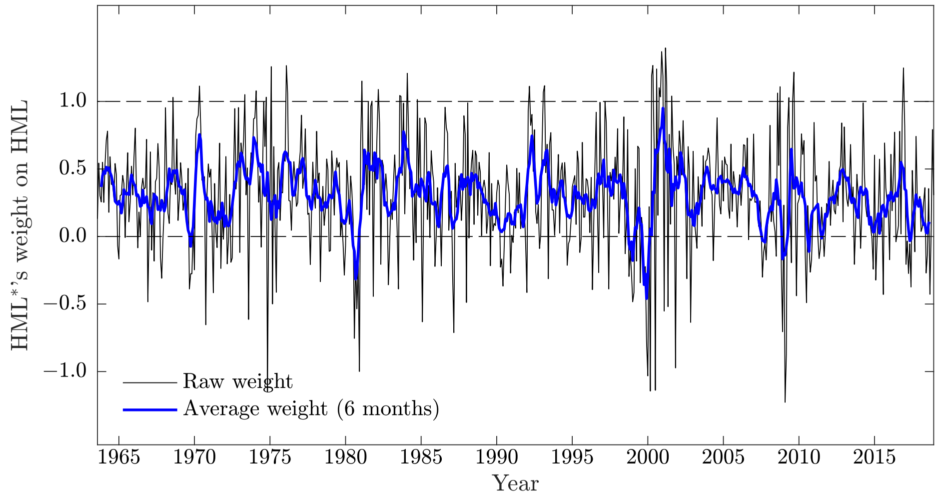

Figure 2: **Efficient HML's weight on standard HML.** This figure shows time-series efficient HML's weight on the standard HML. A time-series efficient factor uses past returns to predict the standard factor's return. The conditioning information in this figure is HML's month $t-1$ return. The investor knows the standard HML's unconditional mean, volatility, and first-order autocorrelation. The black line is the optimal weight from equation (10) with r_{t-1} as the signal; the blue line is the average weight over a one-year window around month t. We set the efficient factor to earn the same expected return as the standard factor. Monthly factor returns begin in August 1963 and end in December 2018.

first fell to zero and then turned negative.

Although the optimal weight can fall below zero, it does in about 20% of the months because of the nonlinear weighting program. We follow Campbell and Thompson (2008) and restrict the weight to positive in our implementation of time-series efficient factors. Table 1 shows that high factor returns predict high returns but low returns do not predict *negative* returns. This empirical regularity resembles Campbell and Thompson's (2008) argument about the conditional equity risk premium should be positive. Campbell and Thompson (2008) show that the restriction on the sign of expected returns is a simple method for avoiding extreme forecasts and that it improves the out-of-sample performance of their predictive regressions. The $[0, +\infty)$ bound also means that the efficient factor maintains a high correlation with the original factor.

In Panel A of Table 2 we report monthly autocorrelations and Sharpe ratios for the original Fama-French factors with the corresponding expected improvement in Sharpe ratios from equa-

Table 2: Time-series efficient five-factor model: Improvements in Sharpe Ratios

Panel A presents monthly Sharpe ratios, autocorrelations, and the corresponding expected efficient-to-original Sharpe ratios for the five Fama-French factors. Panel B compares the predicted and realized Sharpe ratios of the original factors to their time-series efficient counterparts. A time-series efficient factor uses the original factor's month $t-1$ return to predict its month t premium. An efficient factor targets the same expected premium as the original factor. The efficient factor's predicted Sharpe ratio and the z-value for the predicted improvement are computed from equations (12) and (14). We label predictions as "Predicted efficient factors." "Realized efficient factors" are the Sharpe ratios and test statistics of the actual efficient factors. We construct efficient factors using either the full-sample estimates of the original factor's mean, volatility, and first-order autocorrelation ("full sample") or ten-year rolling averages of these inputs ("out of sample"). We use the same parameters for all factors. Last column ("Avg.") is the average of the five factors. The data begin in August 1963 and end in December 2018.

Panel A: Monthly autocorrelations and Sharpe ratios

	Factor					
Statistic	MKTRF	SMB	HML	RMW	CMA	Avg.
Autocorrelation	0.07	0.06	0.17	0.16	0.12	0.11
SR	0.12	0.08	0.12	0.12	0.14	0.11
SR^*/SR (Eq. 13)	1.17	1.26	1.80	1.72	1.33	1.42

Panel B: Predictions vs. realizations for annualized Sharpe ratios

Factor	Estimation		Factor					
definition	method	Statistic	MKTRF	SMB	HML	RMW	CMA	Avg.
Original factor	—	Sharpe ratio	0.41	0.27	0.40	0.40	0.49	0.39
Predicted	Equation	Sharpe ratio	0.48	0.34	0.73	0.69	0.65	0.56
efficient	(13)	ΔSharpe ratio	0.07	0.07	0.32	0.29	0.16	0.17
factor		z-value	0.99	0.83	2.54	2.34	1.74	—
Realized	Full	Sharpe ratio	0.52	0.38	0.62	0.64	0.62	0.56
efficient	sample	ΔSharpe ratio	0.12	0.10	0.22	0.24	0.12	0.16
factor		z-value	1.62	1.39	2.29	2.30	1.67	3.11
	Out	Sharpe ratio	0.45	0.43	0.66	0.63	0.59	0.55
	of	ΔSharpe ratio	0.04	0.16	0.26	0.23	0.09	0.16
	sample	z-value	0.41	1.78	3.13	2.54	1.23	3.22

tion (13). Expected improvement are smallest for the market size and factors because they are the least autocorrelated and the largest for the value and profitability factors. Sharpe ratio and autocorrelation are equal for the average factor, signaling an expected improvement of 42%.

In Panel B we present annualized Sharpe ratios for the original and efficient factors. We first report the *predicted* Sharpe ratios of the time-series efficient factors and the expected z-values for the changes in Sharpe ratios based on the values in Panel A. The predictions suggest that an investor would expect to obtain economically significant gains by timing all five factors. However, considering the sample size of 55 years, Sharpe ratio improvements are likely only statistically significant for HML and RMW factors.

The next rows show the realized Sharpe ratios of time-series efficient Fama-French factors. We estimate factors' means, variances, and serial correlations in this computation using the full sample, and construct the efficient factors using month $t-1$ returns as conditioning information.[9] The results show that all efficient factors indeed outperform the original factors. The realized improvements in Sharpe ratios range from 0.10 (for SMB) to 0.24 (to RMW). The average factor's Sharpe ratio increases from 0.39 to 0.56, an improvement of 41%, as predicted by the analytical analysis. The z-values for improvements also align with the predictions. Before looking at the realized factors, we would predict that the efficient versions of HML and RMW should deliver Sharpe ratio improvements that are significant at the 5% level; the improvement in CMA should be significant at the 10% level; and the improvements in MKTRF and SMB should be statistically insignificant. The data confirm these predictions. The program's inability to produce statistically significant improvements in the market and size factors is not a sign of any shortcomings. Indeed, the program itself correctly suggests, ex ante, that the improvements would likely be statistically undistinguishable from zero.

The computations above build in a lookahead bias by using the factors' full-sample moments. The results, however, are insensitive to this assumption. We construct implementable real-time efficient factors by computing the required moments using ten-year rolling windows. We use the same parameter estimates for all factors by pooling data across factors. The first efficient factors

[9] In Table A1 in the appendix we use factors' average prior one-year returns instead.

are formed in August 1963 based on moments that are estimated from the pooled data of the previous month. We update the moments every month until the estimation window is 10-years after which we use the same fixed window to re-estimate the parameters. The moments converge rapidly because we pool data across all five factors.

Table 2 shows that the Sharpe ratios of all real-time efficient factors also exceed those of the original factors. The gains are small for the market and investment factors but economically and statistically highly significant for the value and profitability factors. The Sharpe ratios of these factors increase by 0.23 and 0.26. The 0.16 increase (z-value = 3.22) in the average factor's Sharpe ratio is close to the increase we get using the full-sample estimates. The reason for the similarity between the full-sample and real-time estimates lies in the stability of the inputs and the robustness of the program. We discuss robustness in Section 5.2 and we show parameter stability in Appendix D.

4.2 Time-series efficient factors in simulations

What do the improvements that the time-series efficient factors display over the original factors tell us about the dynamics of risk and return in the economy? The *predictions* in Table 2 are based on the assumption that conditional variance is constant. The data, in turn, align with these predictions. Is this evidence sufficient to reject processes with time-varying risk in favor of the one with constant risk? We simulate three economies to study the impact of different risk processes on the properties of the resulting time-series efficient factors. We seek to test the hypothesis that an alternative AR(1) process, one in which conditional risk is not constant, may also produce the moments in data like improvements in Fama-French Sharpe ratios.

We assume that the true unobserved autocorrelated factor returns, f, can be estimated with some measurement error:

$$\hat{f}_t = f_t + \epsilon_t \quad (16)$$

$$f_t = \mu + \rho f_{t-1} + e_t \quad (17)$$

where $f_t \sim N(0.32, 2)$ is the return to an autocorrelated latent factor at time t and $\epsilon_t \sim N(0, 2)$ is the noise or measurement error in estimating f_t. For example, $\hat{f}_t$ can be the 30-40-30 portfolio sorted on book-to-market and f_t the true underlying value factor. The researcher only observes $\hat{f}_t \sim N(0.32, \sqrt{8})$ and the total sample $\rho_{\hat{f}_t}$ and implements the Ferson-Siegel procedure.

We consider three economies with different assumptions about conditional risk:

Economy 1: Constant risk	$\text{var}[\epsilon_t \vert f_{t-1}] = \sigma^2$,
Economy 2: Risk increasing in past returns	$\text{var}[\epsilon_t \vert f_{t-1}] = e^{cf_{t-1}}\sigma^2$,
Economy 3: Risk decreasing in past returns	$\text{var}[\epsilon_t \vert f_{t-1}] = e^{-c' f_{t-1}}\sigma^2$,

where $c = 0.4$ and $c' = 0.33$ are scalars chosen to reduce the slope of the exponential function. A value of $c = 0.4$, for example, implies that the volatility of ϵ_t will be 1.9, 3.7, and 7.2 times larger following a one-, two-, or three-standard deviation move in f_{t-1}. Because ϵ_t is independent of the other source of volatility (e_t), it generates a reasonable amount of extra variation in the total volatility of the observed factor. These parameters produce factors $\hat{f}$ whose distributions properties match those of the average Fama-French factor: the average mean, standard deviation, annualized Sharpe ratio, and autocorrelation of the average factor ($\hat{f}$) in all three simulated economies are 0.32%, 2.82%, 0.39, and 0.11.[10]

[10] We have considered values for the scalar c ranging from 0.10 to 1.00. We have also considered alternative functions for modeling conditional risk such as the cumulative normal distribution function, truncated functions wherein risk is increasing in past positive returns, or increasing in past negative returns. The estimates from these alternative simulations are similar to main results—the precise functional form of *how* risk increases or decreases in past returns is not very important.

The assumptions above about the error term's variance do not conflict with time-series models of Engle (1982) and Bollerslev (1986), or with other stochastic-volatility models that describe variance as a function of the *magnitude* of the previous period's error. Neither clustering nor mean reversion in volatility alters the thrust of the simulations provided that these features are independent from the *level* of past returns. We are interested in whether volatility innovations respond to the *level* of past returns rather than the *magnitude* of past returns. The conclusions drawn from the simulations continue to hold when we consider processes in which volatility clusters or reverts independent of the level of expected returns.

We simulate 10,000 factors for 660 months for each economy and then construct their time-series efficient versions using the factors' total sample parameters. The optimal weight in equation (10) assumes constant conditional risk as in the first economy. Because the second and third economies violate this assumption, we refer to the modified factors in these economies as *semi-efficient* factors.

Figure 3 shows the distributions of the improvements in Sharpe ratios in the simulated economies. The vertical lines indicate the actual (out-of-sample) improvements of the Fama-French factors from Table 2. The actual improvements in Fama-French factors resemble those found in the constant risk economy of the middle; they are too high for the increasing-risk economy and too low for the decreasing-risk economy. We apply the Kolmogorov-Sminov test to the null hypothesis that the Sharpe ratio improvements represented by the bars are jointly drawn from a normal distribution with a mean and variance equal to those of the simulated factors. We fail to reject the null for the constant-risk economy with a p-value of 99.1%. The finding that the improvements in the Fama-French factors best match those of the first economy support the hypothesis that (a) conditional risks of the Fama-French factors are constant and that (b) the standard Fama-French factors are unconditionally inefficient.

Panel B shows the distribution of the average improvement across five factors. In every iteration, we simulate five factors, compute their efficient versions and the changes in Sharpe ratios, and average the changes across the five factors. We repeat this 10,000 times. Panel B shows that, by forming efficient factors in the constant risk economy, the average factor's Sharpe ratio increases from 0.39 to 0.55. This 0.16 improvement in the simulations matches the improvement that the average Fama-French factor displays in the actual data in Table 2.

The average factor's Sharpe ratio increases by only 0.03 in the increasing-risk economy. These factors are just as autocorrelated as in the first economy but they do not benefit meaningfully from the program because of the positive trade-off between risk and return. Sharpe ratio improvements are the largest in the decreasing-risk economy in which factors become *less* risky following high returns. This improvement far exceeds that experienced by the actual factors.

The simulation results show that the improvements in the Fama-French factors are consistent with an economy in which factor returns are positively autocorrelated and conditional risk is constant. This conclusion is not sensitive to the specific assumptions of the simulation: any deviation from the constant risk assumption pushes the statistics of the efficient factors and Sharpe ratio improvements away from those of the actual Fama-French factors. These findings suggest that *conditional on past return*, return changes but variance is constant. Fleming et al. (2001) and Moreira and Muir (2017) show that *conditional on past variance*, variance changes but expected return does not. Both of these risk-return dynamics can exist in presence of one another. Factors' means and variances may follow their own, largely independent, time-varying processes.[11]

[11] Moreira and Muir (2017) propose scaling an allocation by the inverse of the factor's variance. Liu et al. (2019) argue that real-time implementation of volatility-scaling is not feasible and Cederburg et al. (2019) show that this scaling does not increase the Sharpe ratio of the typical factor. Goulding and Mazzoleni (2019) solve for the conditions under which Moreira and Muir's (2017) volatility management approach improves Sharpe ratios. While volatility scaling moderates the variance component, it adds volatility to the mean. As a result, although the intercept of the regression of a volatility-managed factor on its original version is positive by construction, the Sharpe ratio can either increase or decrease. It increases if (a) volatility is sufficiently persistent and (b) the factor's unconditional Sharpe ratio low. For example, if the original factor's Sharpe ratio is above one, volatility scaling can never improve Sharpe ratios. If these two conditions are not met, volatility scaling *reduces* the factor's Sharpe ratio.

Panel A. Individual factors

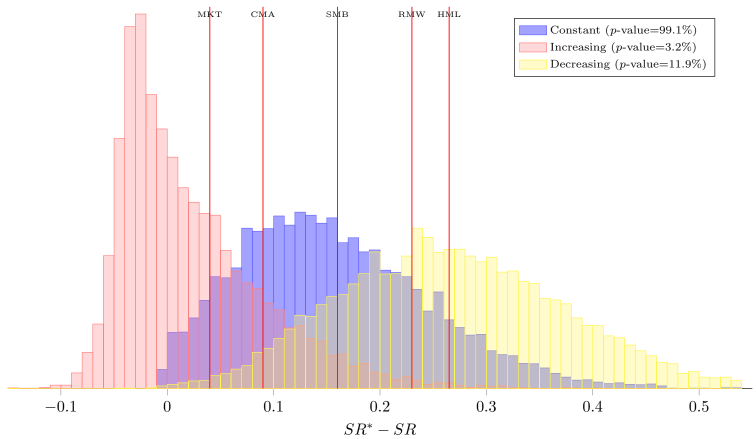

Panel B. Average of five factors

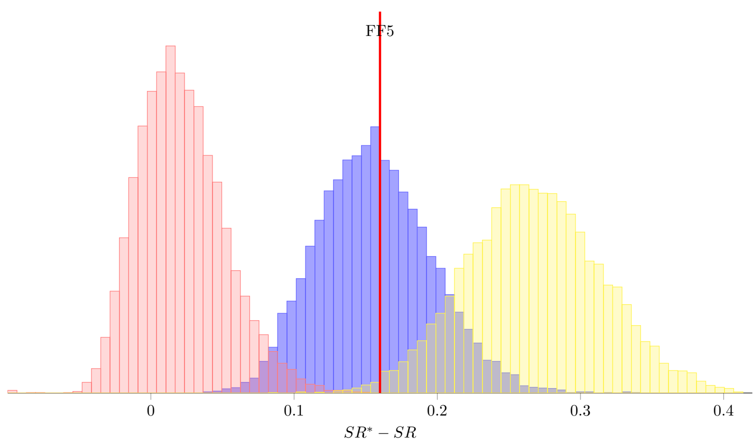

Figure 3: **Simulated improvements in Sharpe ratios: Standard vs. time-series efficient factors.** Panel A shows the distributions of the differences in Sharpe ratios between original and efficient factors. We generate a simulated factor, compute its time-series efficient version, and then compute the difference in Sharpe ratios. We repeat this procedure 10,000 times for each of the three economies that differ in how factor risk responds to past returns. The vertical lines indicate the Sharpe ratios of the five efficient Fama-French factors from Table 2. For each economy we report the p-value from the Kolmogorov-Smirnov test for the null hypothesis that the five actual Fama-French changes in Sharpe ratios are jointly drawn from a normal distribution with a mean and variance equal to the mean and variance of the simulated distribution. Panel B shows the distributions of the changes in the average Sharpe ratio of five randomly selected factors. We simulate five factors, construct the time-series efficient versions, and then compute the difference in the average Sharpe ratios between the efficient and original versions. The vertical line indicate the mean improvement in Sharpe ratios of the five efficient Fama-French factors from Table 2.

Table 3: Time-series efficient five-factor model: Alphas from spanning tests

This table reports alphas and t-values (in parentheses) from regressions in which the dependent variable is one of the factors of the efficient or standard five-factor model and the explanatory variables are all five factors of the other model. The efficient factors are conditional on either the prior-month ($t-1$) or prior-year ($t-12$ to $t-1$) return. We estimate factors' means, volatilities, and first-order autocorrelations each month using a ten-year rolling window. We use the same parameters to construct all efficient factors. The Gibbons et al. (1989) test statistic is under the null hypothesis that the alphas of the five factors are jointly zero. We report p-values associated with these tests in square brackets. Because of the ten-year window, the sample begins in July 1973 and ends in December 2018.

Efficient factors	Dependent variable					GRS test
conditional on:	MKTRF	SMB	HML	RMW	CMA	F-value
			Efficient factors regressed on standard factors			
Prior-month	0.17	0.14	0.17	0.11	0.06	5.77
return	(2.39)	(3.00)	(4.04)	(3.33)	(2.03)	[0.0%]
Prior-year	0.16	0.05	0.12	0.11	0.01	3.35
return	(2.21)	(1.27)	(2.85)	(3.18)	(0.61)	[0.5%]
			Standard factors regressed on efficient factors			
Prior-month	0.30	0.02	0.00	0.07	0.06	2.70
return	(2.49)	(0.29)	(0.04)	(1.29)	(1.51)	[2.0%]
Prior-year	0.16	0.00	0.06	0.05	0.06	1.99
return	(1.80)	(0.02)	(1.13)	(1.00)	(2.23)	[7.9%]

4.3 Time-series efficient versus inefficient factors

4.3.1 Actual factors

Tables 2 and A1 show that efficient factors, constructed using either month $t-1$ returns or average returns from month $t-12$ to $t-1$, earn higher Sharpe ratios than the original factors. In Table 3 we measure the incremental information contents of the two sets of factors relative to each other. We construct the efficient factors in this analysis, and all analyses that follow, using rolling parameters to avoid a lookahead bias. We first regress each of the five efficient factors against all five standard factors and then, reversing the regression, each standard factors against all five efficient factors. We construct the efficient factors using either month $t-1$ returns or average returns from month $t-12$ to $t-1$.

The alphas in Table 3 have two interpretations. First, an insignificant alpha would indicate that the left-hand side factor would not improve the asset pricing model represented by the right-hand side factors (Barillas and Shanken, 2017). For example, when the standard SMB factor's alpha is insignificant in the regression against the time-series efficient five-factor model, this finding suggests that adding the standard SMB factor to the efficient five-factor model would not improve it. Second, from the investment viewpoint, an insignificant alpha indicates that an investor who trades the right-hand side factors could not increase a portfolio's Sharpe ratio by trading also the left-hand side factor (Huberman and Kandel, 1987). The standard SMB factor's insignificant alpha indicates that an investor who trades the efficient five-factor model would not have earned a higher Sharpe ratio by having also traded the standard SMB factor.

The estimates in Table 3 show that the time-series efficient factors typically contain all the information found in the standard factors—and more. The first row shows that time-series efficient factors constructed using month $t-1$ returns are incrementally informative about future returns at the 5% level when controlling for the five-factor model. The Gibbons et al. (1989) test for the joint significance of the alphas returns a p-value < 0.001. In the reverse regressions, except for the market factor, the standard factors are not incrementally informative about future returns when controlling for the efficient five-factor model. However, because of the significance of the market factor's alpha, the GRS test rejects the null that the alphas are jointly zero with a p-value of 0.02.

The differences between original and time-series efficient factors are similar when we condition on prior-year returns. In regressions of the efficient factors against the standard factors, all alphas are positive, and those associated with the market, value, and profitability factors are statistically

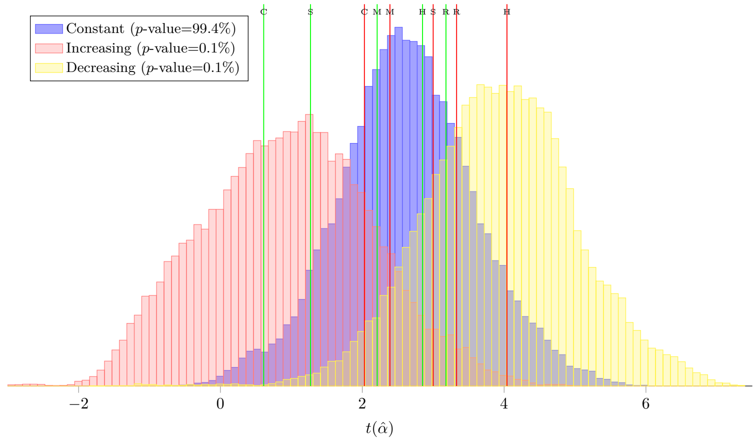

Figure 4: **Simulated alphas: Standard vs. time-series efficient factors.** This figure shows the distributions of $t(\hat{\alpha})$s for time-series efficient factors. We generate five simulated factors, compute their time-series efficient versions, and then regress each efficient factor against all five original factors. We repeat this procedure 10,000 times for each of the three economies that differ in how factor risk responds to past returns. The ten vertical lines indicate the t-values of the ten efficient Fama-French factors from Table 3. The green bars condition on prior one-year returns and the red bars on prior one-month returns. For each economy we report the p-value from the Kolmogorov-Smirnov test for the null hypothesis that the ten actual Fama-French $t(\hat{\alpha})$s are jointly drawn from a normal distribution with a mean and variance equal to the mean and variance of the simulated distribution of $t(\hat{\alpha})$s.

significant at the 1% level. The standard factors, except for the investment factor, are uninformative in the reverse regressions. The GRS test rejects the hypotheses that the alphas are jointly zero with a p-value of 0.005 in the first set of regressions; in the second set of regressions, this p-value is 0.08. Put differently, an investor trading the efficient factors would gain little by trading the standard factors, but an investor trading the standard factors would gain a lot by considering the efficient factors.

4.3.2 Simulated factors

The estimates in Table 3 are also consistent with an economy in which conditional is constant. Figure 4 shows the distribution of $t(\hat{\alpha})$s from spanning regressions that use the same simulated factor data as Figure 3. We estimate regressions in which the dependent variable is one of the efficient factors and the independent variables are all five original factors.

The middle distribution corresponds to the constant-risk economy. The improved factors of this economy are efficient (rather than semi-efficient) because they are computed under the assumption of constant risk. One-quarter of the t-values are less than two. That is, even though the true factors are autocorrelated and homoskedastic, we often fail to reject the null hypothesis that return autocorrelation cannot be transformed into alphas. The vertical bars indicate the t-statistics of the two sets of efficient Fama-French factors from Table 3. The actual t-values fall where most of the simulated density mass is located.The Kolmogorov-Smirnov test does not reject the null hypothesis that the actual $t(\hat{\alpha})$s are drawn from this economy with a p-value of 99.4%.

The distribution of $t(\hat{\alpha})$s shifts to the left in an economy in which risk increases in past returns. Only 20% of these semi-efficient factors earn significant alphas over the original factors. The green and red bars show that the probability of the Fama-French being produced by this return-generating process is low: nine out of the ten intercepts lie to the right of the distribution's mean and only one to its left. The null hypothesis that the alphas are drawn from this distribution is rejected with a p-value $< 0.1\%$.

Semi-efficient factors in the decreasing-economy produce the largest $t(\hat{\alpha})$s. All of the t-values associated with the actual Fama-French factors, except for he HML, lie on the distribution's left side. If factor risk decreased in past returns, improvements should be larger than what they are in the data. We reject the null hypothesis that the ten actual t-values are drawn from this distribution with a p-value of 0.1%.[12] These simulation results, too, suggest that the data-generating process most consistent with the data is the simplest one: factors' expected returns are autocorrelated while risks are constant.

[12] In each simulation we also compute the GRS F-statistic for the null hypothesis that the intercepts of the five (semi-)efficient factors are jointly significant. The distribution of 10,000 simulated GRS F-values of each economy can be compared to the GRS F-values of 5.77 and 3.35 of Table 3. In the first economy the distribution of the F-values has a mean and standard deviation of 7.85 and 2.40. The actual values of 5.77 and 3.35 are therefore within one- and two-standard deviations from the mean of the simulated distribution. The distributions of the GRS F-values for the other two economies are, by contrast, either far to the left (increasing risk) or to the right (decreasing risk) from the actual F-values.

5 Are the results specific to the Fama and French (2015) model?

5.1 Other factor models

The improvements that time-series efficient factors display over the standard factors are not specific to the five-factor model. In Figure 5 we consider the following models:

1. **The Daniel et al. (2020)** hedged-factor model. The factors in this model are the same as those in Fama and French (2015). A hedged factor's return is equal to the original factor's returns minus the return on the hedge portfolio, where each factor's "hedge portfolio" is based on stocks' predicted factor loadings.

2. **The Stambaugh and Yuan (2017) four-factor mispricing model.** The factors in this model are the market and size factors and two factors formed from 11 anomalies of Novy-Marx (2013). The MGMT factor is based on characteristics related to management actions and the PERF factor on those related to firm performance. The size factor deviates from Fama and French in that it only includes stocks not used in forming either of the mispricing factors.

3. **The Hou et al. (2015) *q*-factor model.** The factors in this model are the market (MKT), size (ME), profitability (ROE), and investment (I/A) factors. The market, size, and investment factors are the same as those in the Fama and French (2015) five-factor model; the profitability factor differs in that it is rebalanced *monthly* using return-on-equity from quarterly earnings reports.

4. **Industry-neutral five-factor model.** We follow Novy-Marx (2013) and construct industry-neutral versions of the four non-market factors of the Fama-French model. We first demean each characteristic within the 49 Fama-French industries before sorting stocks into portfolios.

We then overlay an industry hedge: if stock i's weight in a portfolio is w_i, the industry-hedge portfolio invests $-w_i$ in the stock's value-weighted industry portfolio.

5. **The Novy-Marx (2013) four-factor model.** The factors in this model are the market (MKTRF), size (SMB), profitability (PMU), and momentum (UMD) factors. The profitability factor is constructed similar to that in the five-factor model except that, instead of being defined as operating profits-to-book value of equity, Novy-Marx defines the signal as gross profits-to-total assets. The size, profitability, and momentum factors are industry-neutralized using the two-step procedure described above.

We create time-series efficient versions of those factors that are unique to each model. For example, although the Stambaugh and Yuan (2017) and Hou et al. (2015) include the standard market factor, this factor was already in the five-factor model, and we therefore do not add it to Figure 5. We estimate factors' means, volatilities, and first-order autocorrelations each month using a ten-year rolling window using data on all factors shown in this figure and use the same inputs for all efficient factors.

Panel A of Figure 5 shows that in 33 out of 34 comparisons—we have 17 factors, conditioned on either prior-month or prior-year returns—the efficient factor's Sharpe ratio exceeds that of the original factor. Although some improvements are economically and statistically modest, this consistency is marked. The estimates for "All" show that the average average improvement is 0.09 units (from a Sharpe ratio of 0.595 to 0.685) when we condition on prior-month returns and 0.08 when we condition on prior-year returns.

Panel B shows that some of the differences in the Sharpe ratios of the efficient and standard factors are statistically significant in isolation; for the average factor these differences are highly significant. We compute standard errors for the differences for "All" by bootstrapping return data for the efficient and standard factors by calendar month, computing these factors' Sharpe

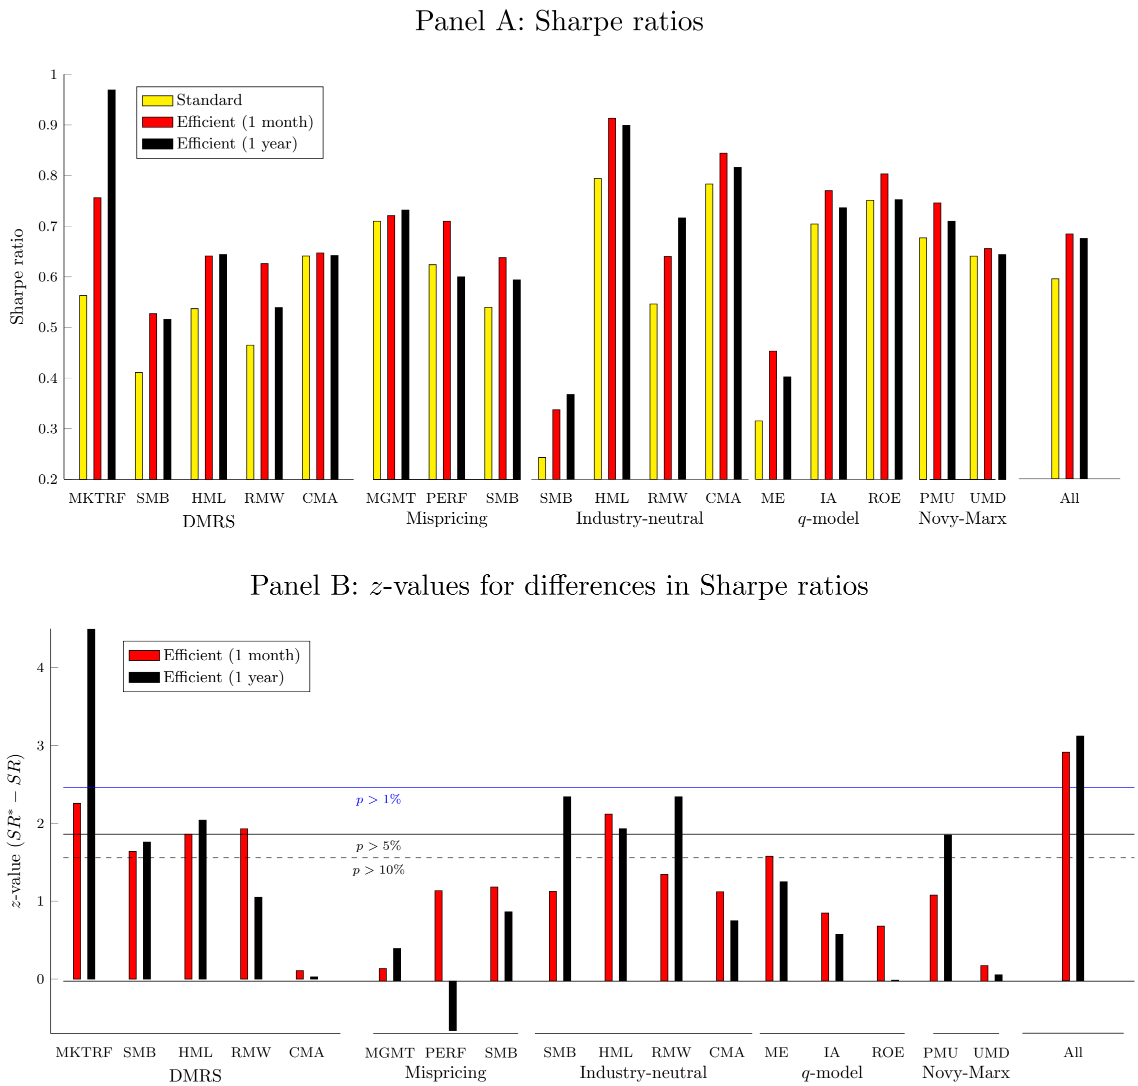

Figure 5: **Other factor models: Standard vs. time-series efficient factors.** Panel A shows the Sharpe ratios of standard factors and their time-series efficient counterparts. Panel B reports z-values for the differences between the efficient and original factors' Sharpe ratios. We construct efficient factors by conditioning on either prior-month (red bars) or prior-year (black bars) returns. The asset pricing models are: (1) Daniel et al. (2020) five-factor model ("DMRS"); (2) Stambaugh and Yuan (2017) four-factor model ("Mispricing"); (3) industry-neutral four-factor model ("Industry-neutral"); (4) Hou et al. (2015) four-factor model (q-model); and (5) Novy-Marx (2013) four-factor model ("Novy-Marx"). We include the factors unique to each model. We estimate factors' means, volatilities, and first-order autocorrelations each month over a ten-year window using data on all factors shown in this figure, and use the same inputs to generate all efficient factors. We compute standard errors for the average factor ("All") by resampling return data for the standard and efficient factors by calendar month. Because of the ten-year window, the sample begins in July 1973 and ends in December 2018, except for the Stambaugh and Yuan (2017) factors, for which the data end in December 2016.

ratios, taking the average of these Sharpe ratios, and then repeat this procedure 10,000 times. The improvement in the average factor's Sharpe ratio has a z-value of 3.05 (prior-month version) or 3.26 (prior-year version).

These results suggest that time-series efficiency, as a method for improving factor's Sharpe ratios, is distinct from and complementary to other improvement protocols. The Daniel et al. (2020) factors, for example, are improved versions of the factors in the five-factor model: they enhance Sharpe ratios by hedging out risk that appears not to be compensated in the cross section of stock returns. Whereas the standard market factor's Sharpe ratio, for example, is 0.41, the hedged market factor has a Sharpe ratio of 0.56. However, because the Daniel et al. (2020) procedure is not concerned with the time-series predictability in factor premiums, time-series efficiency generates additional improvements. The Sharpe ratio of the *time-series efficient* DRMS market factor, for example, is 0.97 when conditioned on prior-year returns, and the improvement over the standard DRMS market factor is significant with a t-value of 4.49, the highest in Figure 5. The gains in Figure 5 result from hedging out unpriced variation in the time-series of factor returns. Similarly, while Novy-Marx's industry neutral HML outperforms the standard HML—these two factors' Sharpe ratios are 0.79 and 0.40—time-series efficiency further increases this factor's Sharpe ratio to 0.90 when conditioned on either prior-year or prior-month returns.

5.2 Robustness: Sensitivity to estimation errors

Time-series efficient factors require estimates of the factors' means, volatilities, and autocorrelations. Table 2 shows that the efficient factors built using these moments' full-sample estimates are similar to those built from moments estimated over ten-year rolling windows. These two sets of efficient factors are similar not only because the parameters are stable but also because the optimal solution is insensitive to estimation errors.

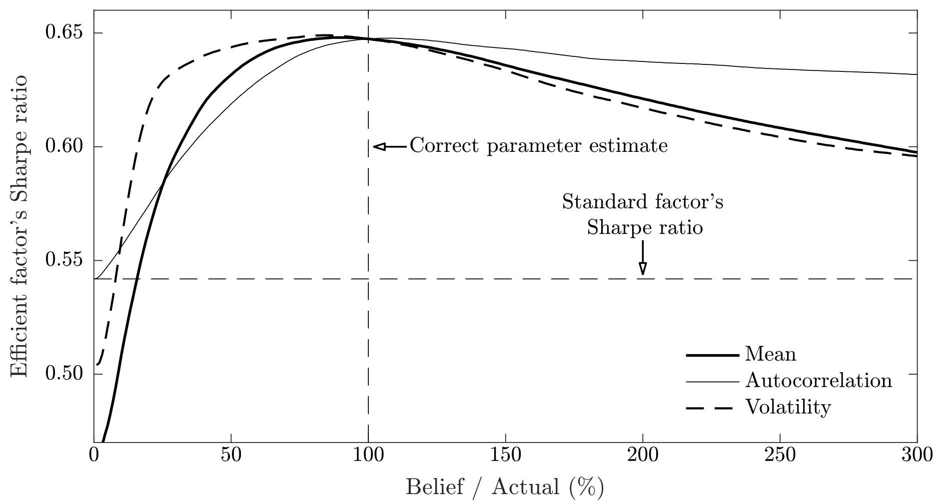

Figure 6: **The sensitivity of the average efficient factor's Sharpe ratio to estimation errors in the inputs.** Time-series efficient factors time the original factors by predicting their month t returns with month $t-1$ returns. The optimal weights depend on the factors' means, volatilities, and first-order autocorrelations. We compute these moments using the full sample from July 1963 through December 2018 for each of the five factors of the Fama and French (2015) model and those listed in Figure 5. The average original factor's Sharpe ratio is 0.54 and the average efficient factor's Sharpe ratio is 0.65. We plot the efficient factor's Sharpe ratio as a function of the investor's beliefs about the inputs. The value on the x-axis is the ratio of the belief to the full-sample parameter estimate; at 100% the investor's belief matches the factor's correct parameter estimate.

Figure 6 shows that investors would typically be better off by constructing the efficient factors even if they suspect that their estimates about the inputs may deviate substantially from the true parameters. In this figure we assume that the investor uses estimates that are between 0% to 300% of the true parameters to generate the efficient factors. The investor uses the correct estimates for the other two parameters. We construct new time-series efficient factors for each set of parameters and compute the Sharpe ratios of the resulting not-quite-as-efficient factors. We repeat this process for every factor and for each of the three inputs and then compute the average Sharpe ratios across all factors for every set of parameters. In this figure we include all factors: the five factors of the Fama-French model and the 17 additional factors from Figure 5. Figure 6 shows how the efficient factors' Sharpe ratios change when we vary the values of the inputs around the full-sample parameter estimates.

The average efficient factor that uses the correct parameters earns a Sharpe ratio of 0.65, which represents a 20% increase over the average standard factor's Sharpe ratio (0.54). Figure 6 shows that even large estimation errors in the inputs do not meaningfully decrease the Sharpe ratios. In terms of autocorrelations, Figure 6 shows that the efficiency gains would evaporate only if the investor would believe that the factors are *not* positively autocorrelated. It would not matter, for example, if the investor were to believe that the autocorrelation is far higher than what it is in the data; the investor would just time the factor more aggressively, holding a bit more of the factor after it has earned positive returns and be more quick to sell it when it has lost money. But, by behaving so, the investor does not lose much relative to the optimal solution computed using the full-sample estimate for the autocorrelation.

An investor would need to underestimate the factor's mean significantly for her to earn a lower Sharpe ratio compared to the buy-and-hold investor. The realized Sharpe ratio peaks when the estimated mean is close to the true mean, but any estimate between 30% to 250% of the true mean results in large improvements. The same is true for the investor's beliefs about volatility. In fact, in terms of realized Sharpe ratios, an investor would have been better off by operating under the assumption that the average factor is somewhat less volatile than what it actually was—an underestimate of $\hat{\sigma} = 0.84 \times \sigma$ gives the very highest Sharpe ratio. Full-sample parameter estimates do not maximize realized Sharpe ratios because the weighting program assumes that factor returns follow AR(1) processes with constant parameter values. The difference between the maximum Sharpe ratio and that based on the full-sample estimates represent the extent to which the data disagree with this assumption.

Table 4: Daily Fama-French five-factor model: Time-series predictability and standard versus efficient factors

Panel A assigns Fama-French factors into terciles based on day $d-1$ returns and reports average day d returns (and t-values) for these terciles. "Average" at the bottom of the table is computed by assigning all factors first into terciles and then computing the average returns for each tercile. Autocorrelation, reported in the rightmost column is between day $d-1$ and d returns. Panel B compares annualized Sharpe ratios of standard time-series efficient factors. A time-series efficient factor uses day $d-1$ return as conditioning information to predict the original factor's return. The efficient factor earns the same expected premium as the original factor; it uses the unconditional mean, volatility, and first-order autocorrelation of the original factor. The data are daily factor returns from July 1963 through December 2018.

Panel A: Daily time-series predictability

	Day $d-1$ return tercile				Auto-
Factor	Low	Middle	High	H−L	correlation
MKTRF	−0.08	0.03	0.12	0.20	0.05
	(−4.46)	(2.62)	(8.73)	(8.96)	(6.08)
SMB	−0.03	0.00	0.05	0.08	0.04
	(−3.75)	(0.17)	(6.43)	(7.24)	(5.30)
HML	−0.06	0.01	0.10	0.16	0.12
	(−7.50)	(2.01)	(11.52)	(13.57)	(14.31)
RMW	−0.03	0.01	0.06	0.09	0.15
	(−5.80)	(3.03)	(9.54)	(11.06)	(17.77)
CMA	−0.04	0.01	0.08	0.12	0.14
	(−7.60)	(1.44)	(12.39)	(14.34)	(17.04)
Average	−0.05	0.01	0.08	0.13	
	(−10.55)	(4.26)	(17.24)	(19.42)	

Panel B: Sharpe ratios of standard and time-series efficient factors

Factor		Factor				
definition	Statistic	MKTRF	SMB	HML	RMW	CMA
Original factor	Sharpe ratio	0.40	0.24	0.51	0.56	0.59
Efficient factor	Sharpe ratio	1.09	0.60	1.12	1.08	1.26
	Δ Sharpe ratio	0.70	0.36	0.61	0.52	0.67
	z-value	6.01	3.10	5.12	4.47	5.86

6 Daily data

The predictability in factor returns is not limited to monthly frequency. In Table 4 we construct time-series efficient Fama-French factors at daily frequency.

Panel A of Table 4 shows that factors returns are significantly positively autocorrelated at the daily frequency. The average daily return on the market factor, for example, varies from −8 basis points to 12 basis points depending on whether this factor's return was in the top or bottom tercile the prior day. The 20-basis point difference in the averages between the top and bottom terciles is significant with a t-value of 8.96. The market factor is not an aberration. The average factor, constructed by sorting all factors independently into terciles based on their day $d-1$ returns, earns an average return of −5 points when in the bottom tercile, a return of 8 basis points when in the top decile, and the difference between the two has a t-value of 19.42. The daily autocorrelations, shown in the rightmost column, are consistent with these tercile sorts. The value, profitability, and investment factors, in particular, are significantly positively autocorrelated. These factors' first-order autocorrelations range from 0.12 to 0.15.

Panel B of Table 4 constructs time-series efficient factors by using this daily predictability in factor premiums to time each factor. Given the amount of predictability, all five efficient versions of the factors significantly outperform their standard counterparts. The annualized Sharpe ratio of the market increases from 0.40 to 1.09; that of the size factor from 0.24 to 0.60; and so forth.[13] The large improvements at the daily frequency are consistent with what equation (13) predicts: because daily and monthly autocorrelations are very similar, the daily Sharpe ratios improve much more because they start at a lower level.

We caution *against* interpreting the results in Table 4 as suggesting that an investor can increase his or her portfolio's Sharpe ratio from 0.40 to 1.09 by switching from the market factor to the time-series efficient market factor. While an investor can earn the Sharpe ratio of 0.40 by following a

[13] The Sharpe ratios of the standard factors in Panel B of Table 4 differ from those reported in Table 2 because of the sampling frequency. In Table 2 a factor's Sharpe ratio is its average monthly return divided by the standard deviation of its monthly returns. In Table 4 it is the factor's average daily return divided by the standard deviations of its *daily* returns. The Sharpe ratios differ because of within-month compounding and because factor returns are autocorrelated; that is, the differences in the Sharpe ratios of the daily and monthly factors in itself indicates that the factors significantly autocorrelate at the daily frequency. A comparison of Sharpe ratios at different frequencies is akin to the variance-ratio test of Lo and MacKinlay (1988).

passive[14] investment strategy, the time-series efficient market factor requires daily rebalancing; the costs of daily rebalancing can negate the much of the "paper" efficiency gains. However, despite any transaction costs, these daily efficient factors serve a vital function: if a researcher benchmarks *new* factors against a model consisting of inefficient factors, these new factors may seem valuable—*but only because they correlate with the inefficiencies found in the original factors.* The daily time-series efficient five-factor model raises the bar for new factors by increasing the squared Sharpe ratio of the ex-post optimal portfolio consisting of the five factors alone to 4.

7 Momentum and the standard and efficient Fama and French (2015) five-factor model

Jegadeesh and Titman (1993) show that a strategy that buys stocks that have earned higher returns relative to other stocks over the past year continue to outperform stocks that have earned relatively low returns. Asness et al. (2013) show that this momentum effect is not specific to just the cross section of equities; the same effect exists also in, for example, fixed income, currency, and commodity markets. Ehsani and Linnainmaa (2019) suggest that momentum is inextricably linked to factors; factors, too, display momentum and factor momentum explains all forms of individual stock momentum.

Time-series efficient factors push factor momentum back into the individual factors. For example, conditioning on the average month $t-12$ to $t-1$ return, the average Fama-French factor's one-year autocorrelation decreases from 0.063 to -0.005. If cross-sectional momentum strategies indeed profit from the time-series predictability found in the original factors, then the time-series efficient factors should explain at least a part of the momentum profits; the point of the efficient

[14] Pedersen (2018) suggests that some caveats apply: an investor who wants to track the market needs to trade as well because the "market" changes with the entry of new and the disappearance of old companies and when existing firms issue new and repurchase old shares.

Table 5: Momentum and the standard versus efficient Fama and French (2015) five-factor model

This table reports estimates from time-series regressions that measure the association between momentum and the standard and efficient versions of the Fama and French (2015) five-factor model. In the first three models the dependent variable is the Carhart (1997) momentum factor. In the other three models, the dependent variable is a time-series momentum strategy that trades the five factors of the Fama-French model. This strategy is long those factors that earned a positive average return over the prior year and short those with negative past returns. The asset pricing model is the CAPM, the standard five-factor model, or the efficient five-factor model. Time-series efficient factors condition on prior-year factor returns. The efficient factor targets the same expected premium as the standard factor; the weight is based on the original factors' unconditional means, volatilities, and first-order autocorrelations. We estimate these inputs using ten-year rolling windows. The sample period begins in July 1973 and ends in December 2018.

	Dependent variable					
	UMD			FF5 Factor Momentum		
Independent		Fama-French FF5			Fama-French FF5	
variable	CAPM	Standard	Efficient	CAPM	Standard	Efficient
Alpha	0.72	0.72	0.16	0.30	0.24	−0.05
	(4.48)	(4.44)	(1.04)	(4.52)	(3.63)	(−0.89)
MKTRF	−0.12	−0.13	0.21	0.02	0.02	0.21
	(−3.33)	(−3.24)	(4.45)	(1.04)	(1.01)	(13.62)
SMB		0.07	0.24		0.10	0.18
		(1.16)	(3.71)		(4.45)	(8.47)
HML		−0.51	−0.02		−0.04	0.21
		(−6.59)	(−0.16)		(−1.33)	(6.49)
RMW		0.24	1.04		−0.02	0.34
		(3.09)	(10.70)		(−0.50)	(10.33)
CMA		0.35	0.35		0.17	0.22
		(3.05)	(2.75)		(3.56)	(5.05)
N	665	665	665	665	665	665
Adjusted R^2	1.5%	8.3%	16.9%	0.0%	5.1%	42.5%

factors is that they largely exhaust the source of such profits.

7.1 Actual factors

In Table 5 we examine the connection between momentum and time-series efficiency. In the first three regressions the dependent variable is the return on the Carhart (1997) momentum factor, UMD. This factor earns a CAPM alpha of 73 basis points (t-value of 4.48) over the 1973–2018 sample period. Fama and French (2016) show that the five-factor model explains none of the momentum profits. Indeed, our second regression shows that UMD's five-factor model alpha is 72 basis points (t-value of 4.44). In the third regression we explain momentum profits using the time-series efficient five-factor model. Because UMD uses past-year returns for sorting stocks into portfolios, we construct efficient factors using the last-year signal and, as before, we estimate the parameters over ten-year rolling windows. The alpha decreases by three quarters to 16 basis points per month, and the t-value to just 1.04.

The standard five-factor model does not price momentum in the time-series regression with constant slopes because these factors are not unconditionally minimum-variance efficient. The nonzero alpha in this regression does not indicate that momentum represents a factor that is unrelated to the model's five factors; rather, it arises because we have not conditioned out the correlation between UMD and factor autocorrelations. The model in the third column is a valid unconditional test of the Fama-French five-factor model because it benchmarks momentum against MVE factors construct using the same past-year return signal.

The dependent variable in the other three columns in Table 5 is a *factor momentum* strategy. This strategy takes long and short positions in the five Fama-French factors based on their average returns from month $t-12$ to $t-1$ and rebalances monthly (Ehsani and Linnainmaa, 2019). If, for example, the market is down but the other four factors are up, then this time-series factor momentum strategy is short the market and long the other four factors. This strategy, which is self-financing because all factors are zero-investment strategies, earns a CAPM alpha of 30 basis points per month (t-value = 4.52). The five-factor model itself cannot explain these profits; its five-factor model alpha is 24 basis points (t-value = 3.63). The six basis points of the alpha that the five-factor model explains relates to the strategy's net long investment in factors (Goyal and Jegadeesh, 2017). The factor momentum strategy is more exposed to factors with higher means, and the five-factor model regression removes this static component of the trading profits.

Because the factor momentum strategy relies on the same feature of returns as time-series efficient factors, this strategy's alpha in the *time-series efficient* five-factor model is close to zero: -5 basis point per month (t-value $= -0.89$). This result implies that an investor who trades the time-series efficient factors would find the factor momentum strategy redundant; the time-series efficient factors already reap the same predictable time variation in factor premiums that the factor momentum strategy targets.

Five-factor model's ability to price momentum is important from the model building perspective. Starting with Fama and French (1996), the dual objective of new asset pricing models has been, first, to capture cross-sectional variation in asset returns and, second, to do so by presenting the most parsimonious model. The ideal is to have a model that does not include a separate factor for each new anomaly discovered in the data. The results in Table 5 represent a step towards parsimony: instead of augmenting the five-factor model with a momentum factor, the original factors can be replaced with their time-series efficient counterparts.

7.2 Simulated factors

We examine the connection between momentum and time-series factor efficiency in the same three simulations as above. In each economy, we simulate five factors, stock betas, and stock residual returns to generate a cross section of returns for 300 stocks for 660 periods. We draw stocks' factor betas from a normal distribution $\beta_F \sim \mathcal{N}(1, 1)$ and keep them constant in each simulation. We draw stocks' residual returns from a normal distribution, $\varepsilon \sim \mathcal{N}(0, 12)$. Stock i's return in month t is

$$r_{i,t} = \sum_{f=1}^{5} \beta_{i,f} F_f + \varepsilon_{i,t}.$$

An individual stock momentum strategy is long an equal-weighted portfolio of the 100 winners

Panel A: Regression results for simulated UMD

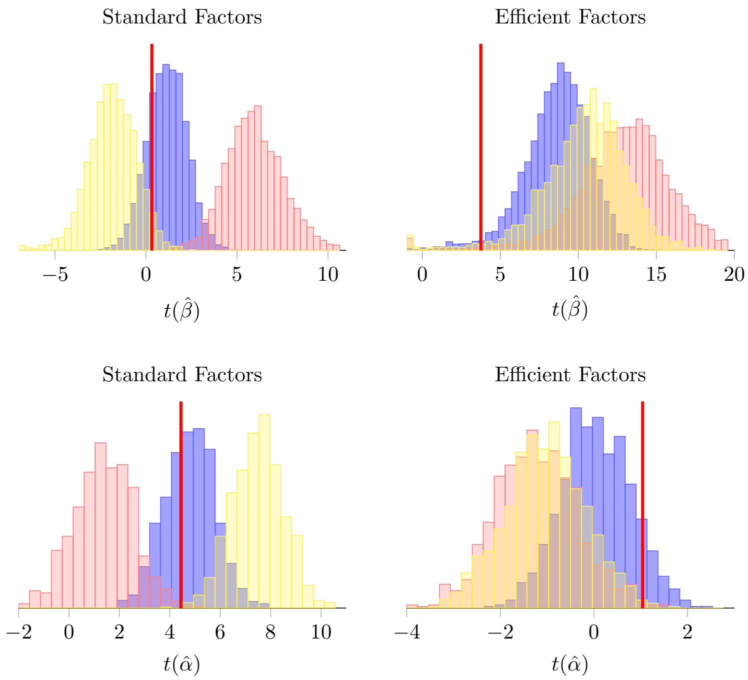

Figure 7: **Momentum, factors, and efficient factors in simulations.** This figure shows the distributions of the slope and intercept estimates from time-series regressions in which the dependent variable is either the simulated individual stock momentum factor (UMD) or the simulated factor momentum strategy (FMOM). We simulate factor returns for constant (blue), increasing (red), and decreasing (yellow) conditional risk economies. In each simulation five factors drive the cross section of stock returns. We construct strategies that trade momentum either in the individual stocks or the five factors. We regress the momentum strategy returns against either the original factors or time-series efficient factors and report $t(\hat{\alpha})$s and $t(\beta_F)$s.

from period $t-1$ and short a portfolio of 100 losers. As factor returns and firm residual returns change every month, firms move from the winner portfolio to the loser portfolio and vice versa.[15] In addition to each simulated UMD we also record the original and efficient (or semi-efficient) five factors. We then regress the simulated UMD or simulated factor momentum on the five factors or their efficient versions. The simulation results should be compared to the numbers in Table 5 which show corresponding results using data.

Panel A of Figure 7 shows the distribution of t-statistics for slopes and intercepts from the

[15]The average return on the UMD in the three economies is 0.71% (constant risk), 0.65% (risk increasing in past returns), and 0.90% (risk decreasing in past returns). The average t-values are 5.66, 4.27, and 6.56.

Panel B: Regression results for simulated Factor Momentum

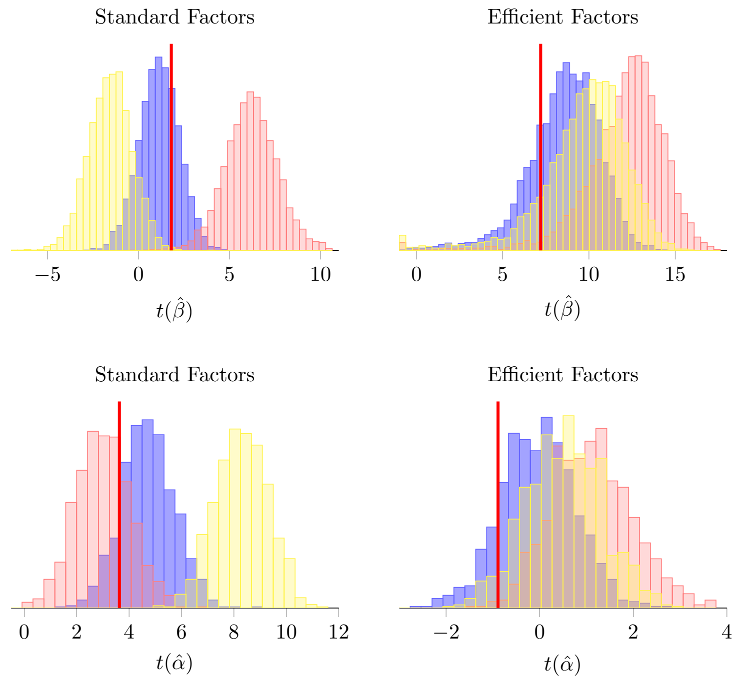

regressions of simulated UMD on the original and efficient factors. The blue, red, and yellow lines represent the constant-, increasing-, and decreasing-risk economies. The vertical line represents the point estimate from the historical data from Table 5.

Simulated UMDs in the constant risk economy do not load significantly on the factors. The $t(\hat{\beta})$s in this economy are small and comparable to those in historical data. As a result, the $t(\hat{\alpha})$s in this economy are also similar to the actual data: the average simulated $t(\hat{\alpha})$ is 4.79; that of the actual UMD is t-value of 4.44 (Table 5). It might seem surprising that an economy governed by five factors (and with no distinct UMD factor) produces momentum returns that *cannot* be described by the same five factors. The impostor "UMD factor" in this economy is an assembly of factor inefficiencies. Because these factor inefficiencies produce the UMD profits, the five factors in an *unconditional* regression cannot price UMD.

There is no momentum anomaly in the increasing risk economy. UMD of this economy loads positively on its factors—$t(\hat{\beta})$s are large—and, therefore, the $t(\hat{\alpha})$s are close to zero. By buying winners and shorting losers, the UMD strategy still bets on return continuation, which exists also in this economy. However, because conditional risk now increases in past returns, the winner factors become more volatile next month and the loser factors less volatile. The UMD factor in economy is therefore also a positive bet on factor *risk*—which is why $t(\hat{\beta})$s for the standard factors are high. Data simulated from this economy are not consistent with the actual data: in this increasing-risk economy, there is no momentum net of the standard factors.

Individual stock momentum is more anomalous in the decreasing-risk economy than in the historical data or the other two simulated economies. The mechanism responsible for this effect is the same as in the second economy except that it now runs in the opposite direction: UMD is implicitly long winner factors and short loser factors, and because risk decreases in past returns, the winner factors are less volatile than loser factors. The momentum strategy therefore now becomes a bet *against* factor risk, which is why the yellow bars of $t(\hat{\beta})$s are mostly negative and significant.[16]

We present the estimates from regressions of UMD on the simulated efficient factors in the right column. UMDs in all economies load significantly on the efficient (or, in the case of the increasing- and decreasing-risk economies, semi-efficient) factors. The loadings are smaller in the constant-risk economy than what they are in the other two economies. That is, deviations from the constant-risk assumption push the distribution away from the actual data.

The distributions of $t(\hat{\alpha})$s indicate that time-series efficient factors explain away momentum profits in all three economies. Time-series efficiency has pushed momentum back into the factor returns. However, the $t(\hat{\alpha})$ of the historical UMD intercept is likely drawn from a constant risk economy. Even a small deviation from the constant-risk assumption pushes the distribution away from the historical data. In all four figures of Panel A the historical red line is closest to the mean of

[16]The mechanism that generates the negative factor loadings in this regression is the same mechanism that produces low-volatility strategies' negative market betas in the actual data. A cross-sectional low-volatility strategy that buys low-volatility stocks and sells high-volatility stocks is not a direct bet against market beta, but it becomes one because the strategy is short market volatility (Ehsani, 2020).

the histogram corresponding to the constant-risk economy. This similarity supports the hypothesis that the autocorrelations in factor returns are unrelated to changes in risk.

In Panel B of Figure 7 we repeat the regressions with the simulated factor momentum strategy as the dependent variable. The conclusions parallel the UMD simulation results. The simulated factors in the constant-risk economy best replicate the slopes and intercepts of the factor momentum in historical data. The $t(\hat{\alpha})$ of the factor momentum strategy are too small in an economy in which risk increases in past returns and too large in the decreasing-risk economy. Factor loadings in a regression of the historical factor momentum strategy on the standard factors are close to zero in the actual date. In the simulations, these slopes are close to zero only in the constant-risk economy.

8 Conclusions

We show that most asset pricing factors are unconditionally inefficient. Following periods of high returns, factors' expected returns are higher but their variances are not. An investor who allocates more to the factors after these periods and scales back after factors have underperformed earns a higher Sharpe ratio than a buy-and-hold investor. We call factors that exhaust their own time-series predictability *time-series efficient factors*.

We show that time-series efficient factors outperform their original counterparts. In the Fama and French (2015) five-factor model, all efficient factors earn higher Sharpe ratios than the original factors, and the time-series efficient factors contain most—or all, depending on the definition of the conditioning information—of the information found in the original factors. This result is not specific to the five-factor model: time-series efficiency improves Sharpe ratios of the factors in all popular asset pricing models. Although some of the improvements in Sharpe ratios are modest, many are both economically and statistically highly significant.

An investor or an asset pricer interested in efficiency does not need to operate in the dark. An investor, given her beliefs about the original factor's moments, can compute *ex ante* the expected efficiency gain. If a factor's return predictability is accompanied by changes in risk or if its unconditional Sharpe ratio is large relative to its autocorrelation, the factor is likely sufficiently time-series efficient, no gains can be expected and the investor need not bother.

What do our results tell us about factors' risk-return dynamics? Throughout our work we compare actual factors to factors simulated from three economies. In these economies factor risk is either constant or increasing or decreasing in past factor returns. The extent to which risk responds to past factor returns determines the extent to which time-series efficiency improves factors' Sharpe ratios. The simplest of these economies—the one in which conditional risk is constant—is the most consistent with all our empirical results. In an economy in which risk increases in past returns, time-series efficiency would not produce nearly as large gains and momentum will not exist. In an economy in which decreasing in past returns, gains would be far greater and a momentum strategy would be even more anomalous. No other assumption has a better chance of replicating historical data than the constant risk.

Our results are important from both investing and asset pricing perspectives. Because factor autocorrelations are persistent, the increases in Sharpe ratios at the monthly frequency can cover any trading costs. The connection between time-series efficiency and the cross-sectional momentum effect buttresses this point: if standard momentum strategies are profitable net of trading costs (Asness et al., 2014), then time-series efficiency must be profitable as well—the two effects stem from the same source.

From the perspective of Cochrane's (2011) factor zoo, time-series efficient factors serve a vital purpose no matter what their transaction costs are. The time-series predictability present in the original factors implies that the original factors are unconditionally inefficient; and that, without introducing additional factors, unconditional tests of the original factor models are invalid. The corollary to this point is that when researchers use the *original* factors as the benchmarks, they may easily mistake indirect efficiency gains for novel anomalies. An additional factor that they propose may seem to yield Sharpe ratio improvements—but only because it correlates with inefficiencies present in the original factors. By using efficient factors in the asset pricing models, we ensure that any proposed factor serves a purpose other than extracting such indirect efficiency gains. For example, because time-series efficiency gains and momentum profits stem from the same source—factor autocorrelations—the momentum factor becomes redundant when an investor or a builder of asset pricing models switches from the original factors to time-series efficient factors.

References

Arnott, Rob, Mark Clements, Vitali Kalesnik, and Juhani T. Linnainmaa, 2019, Factor momentum, Working paper.

Asness, Clifford, Andrea Frazzini, Ronen Israel, and Tobias Moskowitz, 2014, Fact, fiction, and momentum investing, *Journal of Portfolio Management* 40, 75–92.

Asness, Clifford S., Tobias J. Moskowitz, and Lasse Heje Pedersen, 2013, Value and momentum everywhere, *Journal of Finance* 68, 929–985.

Asness, Clifford S., R. Burt Porter, and Ross L. Stevens, 2000, Predicting stock returns using industry-relative firm characteristics, AQR Capital Management working paper.

Barillas, Francisco, and Jay Shanken, 2017, Which alpha? *Review of Financial Studies* 30, 1316–1338.

Bekaert, Geert, and Jun Liu, 2004, Conditioning information and variance bounds on pricing kernels, *Review of Financial Studies* 17, 339–378.

Bollerslev, Tim, 1986, Generalized autoregressive conditional heteroskedasticity, *Journal of econometrics* 31, 307–327.

Campbell, John Y, and Samuel B Thompson, 2008, Predicting excess stock returns out of sample: Can anything beat the historical average?, *The Review of Financial Studies* 21, 1509–1531.

Carhart, Mark M., 1997, On persistence in mutual fund performance, *Journal of Finance* 52, 57–82.

Cederburg, Scott, Michael S. O'Doherty, Feifei Wang, and Xuemin Sterling Yan, 2019, On the performance of volatility-managed portfolios, *Journal of Financial Economics*, forthcoming.

Cochrane, John H, 2011, Presidential address: Discount rates, *Journal of Finance* 66, 1047–1108.

Cohen, Randolph B., and Christopher K. Polk, 1998, The impact of industry factors in asset-pricing tests, Kellogg Graduate School of Management working paper.

Daniel, Kent, Lira Mota, Simon Rottke, and Tano Santos, 2020, The cross section of risk and return, *Review of Financial Studies* 33, 1927–1979.

Dybvig, Philip H, and Jonathan E Ingersoll Jr, 1982, Mean-variance theory in complete markets, *Journal of Business* 233–251.

Ehsani, Sina, 2020, The factor risk in low-risk anomalies, SSRN Working Paper No. 3480257.

Ehsani, Sina, and Juhani Linnainmaa, 2019, Factor momentum and the momentum factor, NBER Working Paper No. 25551.

Engle, Robert F, 1982, Autoregressive conditional heteroscedasticity with estimates of the variance of united kingdom inflation, *Econometrica: Journal of the Econometric Society* 987–1007.

Fama, Eugene F, 2014, Two pillars of asset pricing, *American Economic Review* 104, 1467–85.

Fama, Eugene F, and Kenneth R French, 1996, Multifactor explanations of asset pricing anomalies, *Journal of Finance* 51, 55–84.

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, 2016, Dissecting anomalies with a five-factor model, *Review of Financial Studies* 29, 69–103.

Ferson, Wayne E., and Andrew F. Siegel, 2001, The efficient use of conditioning information in portfolios, *Journal of Finance* 56, 967–982.

Fleming, Jeff, Chris Kirby, and Barbara Ostdiek, 2001, The economic value of volatility timing, *The Journal of Finance* 56, 329–352.

Gibbons, Michael R., Stephen A. Ross, and Jay Shanken, 1989, A test of the efficiency of a given portfolio, *Econometrica* 57, 1121–1152.

Goulding, Christian, and Michele Mazzoleni, 2019, On volatility managed portfolios, Working paper.

Goyal, Amit, and Narasimhan Jegadeesh, 2017, Cross-sectional and time-series tests of return predictability: What is the difference? *Review of Financial Studies* 31, 1784–1824.

Grinblatt, Mark, and Sheridan Titman, 1987, The relation between mean-variance efficiency and arbitrage pricing, *Journal of Business* 97–112.

Hansen, Lars Peter, and Ravi Jagannathan, 1991, Implications of security market data for models of dynamic economies, *Journal of political economy* 99, 225–262.

Hansen, Lars Peter, and Thomas J Sargent, 2008, *Robustness* (Princeton University Press, Princeton, NJ).

Hou, Kewei, Chen Xue, and Lu Zhang, 2015, Digesting anomalies: An investment approach, *Review of Financial Studies* 28, 650–705.

Huberman, Gur, and Shmuel Kandel, 1987, Mean-variance spanning, *Journal of Finance* 42, 873–888.

Huberman, Gur, Shmuel Kandel, and Robert F Stambaugh, 1987, Mimicking portfolios and exact arbitrage pricing, *The Journal of Finance* 42, 1–9.

Jegadeesh, Narasimhan, and Sheridan Titman, 1993, Returns to buying winners and selling losers: Implications for stock market efficiency, *Journal of Finance* 48, 65–91.

Jobson, J. Dave, and Bob M. Korkie, 1981, Performance hypothesis testing with the Sharpe and Treynor measures, *Journal of Finance* 36, 889–908.

Liu, Fang, Xiaoxiao Tang, and Guofu Zhou, 2019, Volatility-managed portfolio: Does it really work? *Journal of Portfolio Management* 46, 38–51.

Lo, Andrew W, and A Craig MacKinlay, 1988, Stock market prices do not follow random walks:

Evidence from a simple specification test, *Review of Financial Studies* 1, 41–66.

Markowitz, Harry, 1959, *Portfolio Selection: Efficient Diversification of Investments* (John Wiley, New York).

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

Memmel, Christoph, 2003, Performance hypothesis testing with the Sharpe ratio, *Finance Letters* 1, 21–23.

Moreira, Alan, and Tyler Muir, 2017, Volatility-managed portfolios, *Journal of Finance* 72, 1611–1644.

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

Pedersen, Lasse Heje, 2018, Sharpening the arithmetic of active management, *Financial Analysts Journal* 74, 21–36.

Stambaugh, Robert F., and Yu Yuan, 2017, Mispricing factors, *Review of Financial Studies* 30, 1270–1315.

Appendices

A Autocorrelation may or may not impact portfolio choice

Figure A1 illustrates the effect of return predictability—or the lack thereof—on portfolio choice. We consider an informed investor with perfect information on the log-risk premiums ($\mathbb{E}(\tilde{r}) - r_f$) and variances ($\sigma^2$) for two consecutive months. The investor's objective is to choose the optimal period 2 allocation to maximize the total Sharpe ratio over the two periods.

In Scenario A we choose the return and variance parameters to maintain a constant Sharpe ratio in both periods. A buy-and-hold investor would earn a total risk premium of $0.01 + 0.02 = 0.03$ at a variance of $0.01 + 0.04 = 0.05$, which gives a Sharpe ratio of 0.13. The informed investor in this scenario knows that the second period's expected return is higher. However, she finds it optimal to decrease the allocation when the expected return is higher because so is variance. The investor earns the maximum Sharpe ratio of 0.14 by decreasing the allocation from 100% to 50%. If, instead, the investor were to increase the allocation in the second period to 200%, her total Sharpe ratio would be just 0.12. Even if the second-period Sharpe ratio was slightly higher than the first-period Sharpe ratio—for example, let $\mathbb{E}(\tilde{r}_2) - r_f = 0.02$ and $\sigma_2^2 = 0.03$ for a second-period Sharpe ratio of 0.12—the investor would still find it optimal to decrease the allocation.

Scenario B keeps the variance constant. A buy-and-hold investor now earns a Sharpe ratio of $\frac{0.03}{\sqrt{0.02}} = 0.21$. An investor who knows that the expected return in the second period is higher, while the variance is not, increases her allocation from 100% to 200% in the second period. By doing so, she earns the maximum Sharpe ratio of 0.22. If variances increase *proportional* to the changes in expected returns, an investor maximizing Sharpe ratios does not try to time the asset. However, if risks do not change by as much as expected returns, an investor should increase the allocation as the expected return increases. These two scenarios illustrate that the solution to the allocation

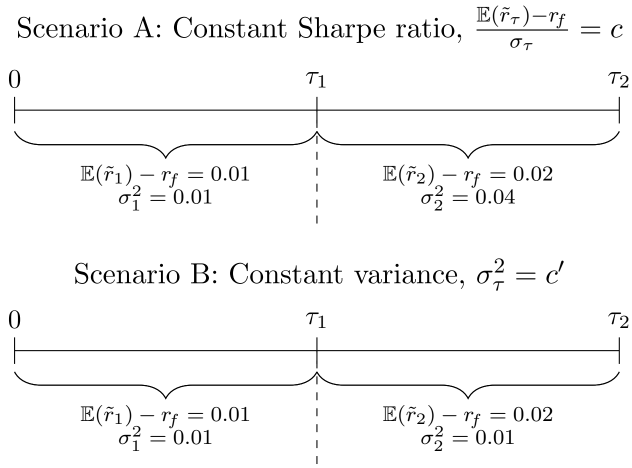

Figure A1: **Return autocorrelation may not affect portfolio choice.** This figure shows hypothetical returns and variances of two risky assets over two periods. The parameters in the two scenarios are identical except for the second-period variance.

problem depends on how risks change when expected returns change. Timing the autocorrelated factor benefits the investor only when the changes in risks are not proportional to the changes in expected returns.

B Deriving ζ and the optimal-to-original Sharpe ratio in equation (13)

Constant ζ is required for computing the optimal weight in the Ferson-Siegel procedure. In this appendix we derive ζ of the weight function of the portfolio when the portfolio's past return is used to forecast future returns: $\tilde{R} = \mu(\tilde{S}) + \tilde{\epsilon}$, in which $\mu(\tilde{S}) = \mu + \rho\tilde{R}_{t-1}$. The portfolio's unconditional expected return is $\mu_p = \frac{\mu}{1-\rho}$. This portfolio's Sharpe ratio is

$$SR = \frac{\mu_p}{\sqrt{\sigma^2_{\mu(\tilde{S})} + \sigma^2_{\epsilon(\tilde{S})}}}. \tag{A-1}$$

Because returns follow an AR(1) process, we can write

$$\frac{\sigma^2_{\mu(\tilde{S})}}{\sigma^2_{\mu(\tilde{S})} + \sigma^2_{\epsilon(\tilde{S})}} = \rho^2. \tag{A-2}$$

This ratio can be interpreted as the R^2 from a regression of returns on past returns. In a univariate regression R^2 is the squared the correlation coefficient, which in this case is also the AR(1) coefficient, ρ. Rearranging and assuming homoskedasticity, $\sigma^2_\epsilon(\tilde{S}) = \sigma^2_\epsilon$, we have

$$\sigma^2_\epsilon = \left(\frac{1}{\rho^2} - 1\right)\sigma^2_{\mu(\tilde{S})}. \tag{A-3}$$

We can now compute ζ as follows:

$$\begin{aligned}
\zeta &= \mathbb{E}\left[\frac{\mu^2(\tilde{S})}{\mu^2(\tilde{S})+\sigma_\epsilon^2}\right] = \mathbb{E}\left[1-\frac{\sigma_\epsilon^2}{\mu^2(\tilde{S})+\sigma_\epsilon^2}\right] = 1-\sigma_\epsilon^2\,\mathbb{E}\left[\frac{1}{\mu^2(\tilde{S})+\sigma_\epsilon^2}\right] \\
&= 1-\frac{\sigma_\epsilon^2}{\mathbb{E}\big[\mu^2(\tilde{S})+\sigma_\epsilon^2\big]}+\frac{\mathrm{var}[\mu^2(\tilde{S})+\sigma_\epsilon^2]}{\mathbb{E}[\mu^2(\tilde{S})+\sigma_\epsilon^2]^3} \approx 1-\frac{\sigma_\epsilon^2}{\mathbb{E}\big[\mu^2(\tilde{S})\big]+\sigma_\epsilon^2} \\
&= 1-\frac{\sigma_\epsilon^2}{\mu^2+\sigma_{\mu(S)}^2+\sigma_\epsilon^2} = 1-\frac{\sigma_\epsilon^2}{SR^2[\sigma_{\mu(S)}^2+\sigma_\epsilon^2]+\sigma_{\mu(S)}^2+\sigma_\epsilon^2} \\
&= \frac{SR^2+\rho^2}{SR^2+1}.
\end{aligned} \tag{A-4}$$

The approximation on the second line uses Taylor series expansion and homoskedasticity in returns, and the final result uses equation (A-3) to remove variances from the identity.[17]

We now derive equation (13). The mean and variance of the Ferson-Siegel portfolio are μ_p and $\mu_p^2(\frac{1}{\zeta}-1)$. The Sharpe ratio squared becomes:

$$SR^{*2} = \frac{\mu_p^2}{\mu_p^2(\frac{1}{\zeta}-1)} = \frac{1}{\frac{1}{\zeta}-1} = \frac{\zeta}{1-\zeta} = \frac{\frac{SR^2+\rho^2}{SR^2+1}}{1-\frac{SR^2+\rho^2}{SR^2+1}} = \frac{SR^2+\rho^2}{1-\rho^2} = \frac{SR^2(1+(\frac{\rho}{SR})^2)}{1-\rho^2}. \tag{A-5}$$

We can therefore write the ratio of the efficient factor's Sharpe ratio to the original factor's Sharpe ratio as

$$\frac{SR^*}{SR} = \sqrt{\frac{1+(\frac{\rho}{SR})^2}{1-\rho^2}}, \tag{A-6}$$

which is equation (13).

[17] Using the Taylor series expansion around $\mathbb{E}[X]$ we have:

$$\mathbb{E}\left[\frac{1}{X}\right] \approx \mathbb{E}\left[\frac{1}{\mathbb{E}[X]}-\frac{1}{\mathbb{E}[X]^2}\big(X-\mathbb{E}[X]\big)+\frac{1}{\mathbb{E}[X]^3}\big(X-\mathbb{E}[X]\big)^2\right] = \frac{1}{\mathbb{E}[X]}+\frac{\mathrm{var}[X]}{\mathbb{E}[X]^3}$$

In our framework, the variance of X is the "vol of vol," which is negligible under the homoskedasticity assumption.

C Covariance between the efficient and original factors

Equations (14) and (15) give the test statistic z for the expected difference in the Sharpe ratios between the efficient and original factors. In addition to the means and standard deviations, this test statistic also depends on the covariance between the efficient and original factors. We derive an approximation for this covariance in this appendix.

Omitting subscripts, we let R denote the return to the original factor and x the efficient factor's weight on the original factor. The return on the efficient factor is thus xR and the covariance that we need to compute is $\mathrm{cov}(xR, R)$. This covariance can be expressed as:

$$\begin{aligned}
\mathrm{cov}(xR, R) &= \mathbb{E}[xR^2] - \mathbb{E}[xR]\mathbb{E}[R] \\
&= \mathbb{E}[xR^2] - \mathbb{E}[R]\big[\mathrm{cov}(x, R) + \mathbb{E}[x]\mathbb{E}[R]\big] \\
&= \mathbb{E}[xR^2] - \mathbb{E}[R]\mathrm{cov}(x, R) - \mathbb{E}[R]^2\mathbb{E}[x].
\end{aligned} \tag{A-7}$$

We therefore have to compute $\mathbb{E}[R] = \mu_p$, $\mathbb{E}[x]$, $\mathrm{cov}(x, R)$, and $\mathbb{E}[xR^2]$. We start by computing $\mathbb{E}[x]$:

$$\begin{aligned}
\mathbb{E}[x] &= \frac{\mu_p}{\zeta}\,\mathbb{E}\left[\frac{\mu(\tilde{S})}{\mu(\tilde{S})^2 + \sigma_\epsilon^2}\right] \\
&= \frac{\mu_p}{\zeta}\left[\mathbb{E}[\mu(\tilde{S})] \cdot \mathbb{E}[\frac{1}{\mu(\tilde{S})^2 + \sigma_\epsilon^2}] + \mathrm{cov}\big(\mu(\tilde{S}), \frac{1}{\mu(\tilde{S})^2 + \sigma_\epsilon^2}\big)\right] \\
&\approx \frac{\mu_p}{\zeta}\left[\frac{\mu_p}{\mu^2 + \sigma_{\mu(S)}^2 + \sigma_\epsilon^2}\right] = \frac{\mu_p}{\zeta}\left[\frac{\mu_p}{(\sigma_\mu^2 + \sigma_\epsilon^2)(1 + SR^2)}\right] \\
&= \frac{SR^2}{1 + SR^2} \cdot \frac{1}{\zeta} = \frac{SR^2}{1 + SR^2} \cdot \frac{SR^2 + 1}{SR^2 + \rho^2} = \frac{SR^2}{SR^2 + \rho^2}.
\end{aligned} \tag{A-8}$$

Equation (A-8) suggests that the efficient factor's average time-series exposure to a factor decreases as the factor's autocorrelation coefficient increases. In other words, the strategy "trusts" the signal (past returns) more than the unconditional factor's Sharpe ratio when the autocorrelation coefficient is high. By contrast, when the autocorrelation is low, the strategy prefers to invest less in the signal and more in the factor to secure the unconditional Sharpe ratio provided by the factor.

We next compute $\mathrm{cov}(x, R)$:

$$\begin{aligned}
\mathrm{cov}(x, R) &= \mathrm{cov}\left(\frac{\mu_p}{\zeta}\frac{\mu(\tilde{S})}{\mu(\tilde{S})^2+\sigma_\epsilon^2}\,,\,\mu(\tilde{S})+\tilde{\epsilon}\right) = \mathrm{cov}\left(\frac{\mu_p}{\zeta}\frac{\mu+\rho\tilde{R}_{t-1}}{\mu(\tilde{S})^2+\sigma_\epsilon^2}\,,\,\mu+\rho\tilde{R}_{t-1}+\tilde{\epsilon}\right) \\
&= \frac{\mu_p}{\zeta}\mathrm{cov}\left(\frac{\rho\tilde{R}_{t-1}}{\mu(\tilde{S})^2+\sigma_\epsilon^2}\,,\,\rho\tilde{R}_{t-1}\right) = \frac{\mu_p}{\zeta}\rho^2(\sigma^2_{\mu(\tilde{S})}+\sigma_\epsilon^2)\cdot\mathbb{E}[\frac{1}{\mu(\tilde{S})^2+\sigma_\epsilon^2}] \\
&= \frac{\mu_p}{\zeta}\rho^2(\sigma^2_{\mu(\tilde{S})}+\sigma_\epsilon^2)\cdot[\frac{1}{(\sigma^2_{\mu(\tilde{S})}+\sigma_\epsilon^2)(1+SR^2)}] = \frac{\mu_p}{\zeta}\rho^2\frac{1}{(1+SR^2)} \\
&= \mu_p\cdot\frac{SR^2+1}{SR^2+\rho^2}\cdot\rho^2\cdot\frac{1}{(1+SR^2)} = \frac{\mu_p\cdot\rho^2}{SR^2+\rho^2}.
\end{aligned} \tag{A-9}$$

The term $\mathbb{E}[xR^2]$ can be computed as follows:

$$\begin{aligned}
\mathbb{E}[xR^2] &= \mathbb{E}[x]\cdot\mathbb{E}[R^2]+\mathrm{cov}(x,R^2) = \frac{SR^2}{SR^2+\rho^2}\cdot(\mu_p^2+\sigma^2_{\mu(\tilde{S})}+\sigma^2_{\epsilon(\tilde{S})})+0 \\
&= \frac{SR^2}{SR^2+\rho^2}\cdot\left(\mu_p^2+\left(\frac{\mu_p}{SR}\right)^2\right) = \frac{SR^2}{SR^2+\rho^2}\cdot\mu_p^2\left(1+\frac{1}{SR^2}\right) \\
&= \frac{SR^2}{SR^2+\rho^2}\cdot\mu_p^2\left(\frac{1+SR^2}{SR^2}\right) = \mu_p^2\cdot\frac{SR^2+1}{SR^2+\rho^2}.
\end{aligned} \tag{A-10}$$

Putting it all together, we can compute $\mathrm{cov}(xR, R)$ as follows:

$$\begin{aligned}
\mathrm{cov}(xR, R) &= \mathbb{E}[xR^2]-\mathbb{E}[R]\mathrm{cov}(x,R)-\mathbb{E}[R]^2\mathbb{E}[x] \\
&= \mu_p^2\cdot\frac{SR^2+1}{SR^2+\rho^2}-\mu_p\cdot\frac{\mu_p\cdot\rho^2}{SR^2+\rho^2}-\mu_p^2\cdot\frac{SR^2}{SR^2+\rho^2} \\
&= \frac{\mu_p^2(SR^2+1-\rho^2-SR^2)}{SR^2+\rho^2} = \frac{\mu_p^2(1-\rho^2)}{SR^2+\rho^2}.
\end{aligned} \tag{A-11}$$

We use this covariance between the efficient and original factors to compute the z-value for the *predicted* improvement in Sharpe ratios.

D Means, standard deviations, and first-order autocorrelations of the Fama-French factors

A time-series efficient factor uses a factor's past return together with three input parameters to time the original factors. These three parameters are the original factor's mean, standard deviation, and first-order autocorrelation. Figure A2 shows rolling averages of these inputs for the five factors of the Fama-French model. At the end of every month starting in June 1973, we compute these moments for all five factors using ten years of historical data. We then average these moments over the five factors.

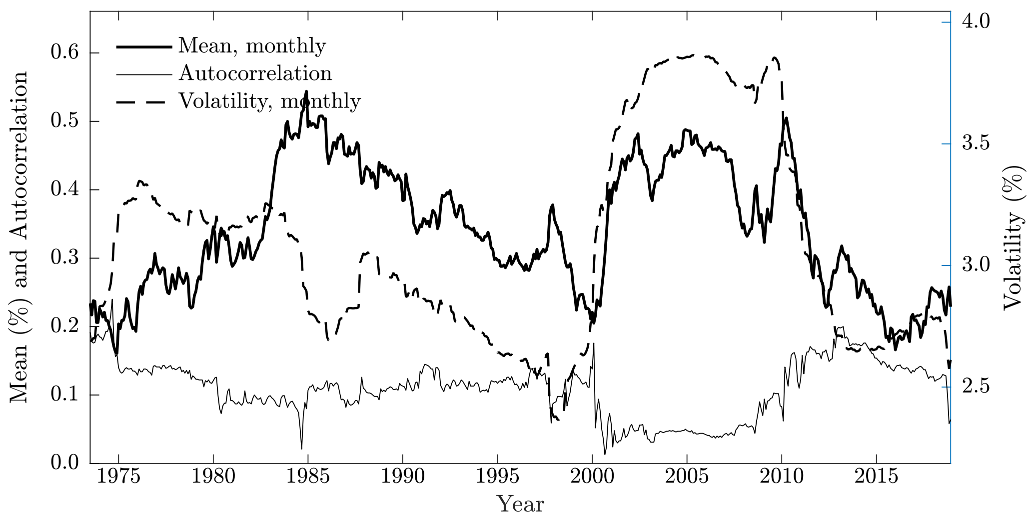

Figure A2: **Means, volatilities, and autocorrelations of the Fama and French (2015) factors.** This figure shows ten-year rolling average estimates of means, standard deviations, and first-order autocorrelations for the five factors of the Fama and French (2015) model. We estimate these parameters separately for each factor and then average over the five factors. Means and standard deviations are in percentages per month. The data are monthly factor returns from July 1963 through December 2018.

Although factors' autocorrelations, for example, vary over time, they remain positive over all ten-year periods. Up to 1997, the average factor's first-order autocorrelation was typically above 0.1; and after a period of lower (but positive) autocorrelations, they again rose above 0.1 in 2010. The estimates in Figure A2 suggest that investor could have in real time made precise inferences about the required moments from historical data alone. The estimates in Table 2, of course, confirm this conjecture: efficient factors based on real-time data perform as well as those based on the full-sample estimates.

Table A1: Fama-French five-factor model: Predictability and efficient factors conditional prior one-year returns

In Panel A we assign factors into two groups based on their returns from month $t-12$ to $t-1$ and report average annualized returns and volatilities for each group. "Average" at the bottom of the table averages across the factors in the high and low groups. In Panel B we compare the original factors of the Fama-French five-factor model to their time-series efficient counterparts. A time-series efficient factor uses factors' average returns from month $t-12$ to $t-1$ to predict the original factor's month t premium. The efficient factor targets the same expected premium as the original factor. The analyses in this table are the same those in Tables 1 and 2 except for the use of the prior one-year return as the conditioning information. Monthly factor returns begin in August 1963 and end in December 2018 except for the rolling estimation, in which the returns begin in July 1973.

Panel A: Returns and volatilities conditional on average returns from month $t-12$ to $t-1$

	Predicting returns			Predicting volatility		
	Prior-year return			Prior-year return		H/L variance
Factor	Low	High	H−L	Low	High	ratio
Market	6.00	6.34	0.34	16.59	13.73	0.68
	(1.90)	(2.43)	(0.08)			[0.00]
Size	−1.26	7.01	8.26	9.95	10.87	1.19
	(−0.67)	(3.39)	(2.95)			[0.11]
Value	1.33	6.51	5.18	10.16	9.18	0.82
	(0.69)	(3.73)	(1.99)			[0.06]
Profitability	1.66	4.30	2.64	9.18	5.27	0.70
	(1.08)	(3.33)	(1.31)			[0.00]
Investment	1.79	5.03	3.24	8.10	6.78	1.21
	(1.43)	(3.66)	(1.74)			[0.09]
Average	1.28	6.17	4.89	7.85	7.65	0.95
	(1.16)	(6.39)	(3.34)			[0.51]

Panel B: Time-series efficient factors

Factor	Estimation		Factor					
definition	method	Statistic	MKTRF	SMB	HML	RMW	CMA	Avg.
Original factor	—	Sharpe ratio	0.41	0.27	0.40	0.40	0.49	0.37
Realized	Full	Sharpe ratio	0.50	0.35	0.48	0.48	0.52	0.44
efficient	sample	ΔSharpe ratio	0.10	0.07	0.08	0.08	0.02	0.07
factor		z-value	1.51	1.84	1.66	1.35	0.99	1.14
	Out	Sharpe ratio	0.52	0.31	0.50	0.57	0.49	0.45
	of	ΔSharpe ratio	0.12	0.03	0.09	0.17	0.00	0.08
	Sample	z-value	1.54	0.60	1.38	2.07	−0.10	2.06