

Is There a Replication Crisis in Finance?

Theis Ingerslev Jensen, Bryan Kelly, and Lasse Heje Pedersen*

First version: August 12 2020. This version: February 10, 2022

Abstract

Several papers argue that financial economics faces a replication crisis because the majority of studies cannot be replicated or are the result of multiple testing of too many factors. We develop and estimate a Bayesian model of factor replication, which leads to different conclusions. The majority of asset pricing factors: (1) can be replicated, (2) can be clustered into 13 themes, the majority of which are significant parts of the tangency portfolio, (3) work out-of-sample in a new large data set covering 93 countries, and (4) have evidence that is strengthened (not weakened) by the large number of observed factors.

Keywords: asset pricing, factors, data mining, replication, multiple testing, external validity, empirical Bayes, Bayesian statistics

JEL Codes: G11, G12, G14, G15, G4, C5

*Jensen is at Copenhagen Business School. Kelly is at Yale School of Management, AQR Capital Management, and NBER; www.bryankellyacademic.org. Pedersen is at AQR Capital Management, Copenhagen Business School, and CEPR; www.lhpedersen.com. We are grateful for helpful comments from Nick Barberis, Andrea Frazzini, Cam Harvey (discussant), Antti Ilmanen, Ronen Israel, Andrew Karolyi, John Liew, Toby Moskowitz, Stefan Nagel, Scott Richardson, Anders Rønn-Nielsen, Neil Shephard (discussant), and seminar and conference participants at AFA 2022, NBER 2021, AQR, Georgetown Virtual Fintech Seminar, Tisvildeleje Summer Workshop 2020, Yale, and the CFA Institute European Investment Conference 2020. We thank Tyler Gwinn for excellent research assistance. Jensen and Pedersen gratefully acknowledge support from the FRIC Center for Financial Frictions (grant no. DNRF102). AQR Capital Management is a global investment management firm, which may or may not apply similar investment techniques or methods of analysis as described herein. The views expressed here are those of the authors and not necessarily those of AQR.

Several research fields face replication crises (or credibility crises), including medicine (Ioannidis, 2005), psychology (Nosek et al., 2012), management (Bettis, 2012), experimental economics (Maniadis et al., 2017), and now also financial economics. Challenges to the replicability of finance research take two basic forms:

1. **No internal validity.** Most studies cannot be replicated with the same data (e.g., because of coding errors or faulty statistics) or are not robust in the sense that the main results cannot be replicated using slightly different methodologies and/or slightly different data.¹ E.g., Hou et al. (2020) state:

“Most anomalies fail to hold up to currently acceptable standards for empirical finance.”

2. **No external validity.** Most studies may be robustly replicated, but are spurious and driven by “*p*-hacking,” that is, finding significant results by testing multiple hypotheses without controlling the false discovery rate. Such spurious results are not expected to replicate in other samples or time periods, in part because the sheer number of factors is simply too large, and too fast growing, to be believable. E.g., Cochrane (2011) asks for a consolidation of the “factor zoo,” and Harvey et al. (2016) states:

*“most claimed research findings in financial economics are likely false.”*²

We examine both of these challenges theoretically and empirically. We conclude that neither criticism is tenable. The majority of factors do replicate, do survive joint modeling of all factors, do hold up out-of-sample, are strengthened (not weakened) by the large number of observed factors, are further strengthened by global evidence, and the number of factors can be understood as multiple versions of a smaller number of themes.

¹Hamermesh (2007) contrasts “pure replication” and “scientific replication.” Pure replication is, “checking on others’ published papers using their data,” also called “reproduction” by Welch (2019). Scientific replication uses, “different sample, different population and perhaps similar, but not identical model.” We focus on scientific replication. We propose a new modeling framework to jointly estimate factor alphas, we use robust factor construction methods that are applied uniformly to all factors, and we test both internal and external validity of prior factor research in several dimensions, including out-of-sample time series replication and international sample replication. In complementary and contemporaneous work, Chen and Zimmermann (2020a) consider pure replication, attempting to use the same data and methods as the original papers for a large number of factors. They are able to reproduce nearly 100% of factors, but Hou et al. (2020) challenge the scientific replication and Harvey et al. (2016) challenge validity due to multiple testing.

²Similarly, Linnainmaa and Roberts (2018) state “the majority of accounting-based return anomalies, including investment, are most likely an artifact of data snooping.”

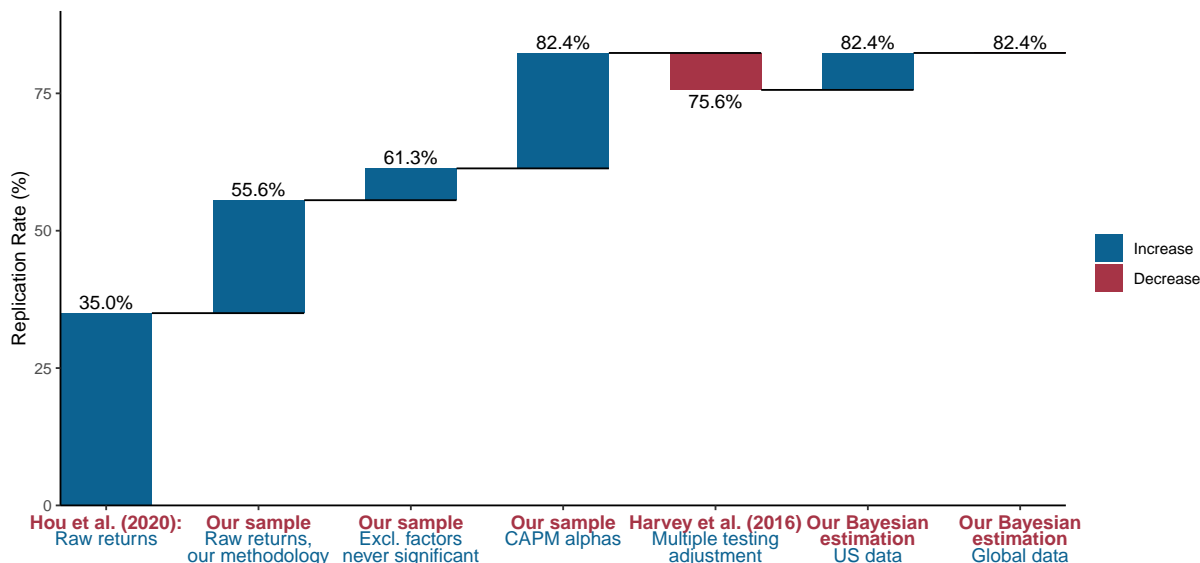


Figure 1: Replication Rates Versus the Literature

Note: This figure summarizes analyses throughout the paper. Refer to Section 3 for estimation details.

These conclusions rely on new theory and data: First, we show that factors must be understood in light of economic theory and we develop a Bayesian model that offers a very different interpretation of the evidence on factor replication. Second, we put together a new global data set of 153 factors across 93 countries. To help advance replication in finance, we have made this data set easily accessible to researchers via a direct open-source link to WRDS, including meticulous documentation of the data and the underlying code base.

Replication results. Figure 1 illustrates our main results and how they relate to the literature in a sequence of steps. It presents the “replication rate,” that is, the percent of factors with a statistically significant average excess return. The starting point of Figure 1—shown as the first bar on the left—is the 35% replication rate reported in the expansive factor replication study of Hou et al. (2020).

The second bar in Figure 1 shows a 55.6% baseline replication rate in our main sample of US factors. It is based on significant OLS t -statistics for average raw factor returns, in direct comparability to the 35% calculation from Hou et al. (2020). This difference arises because our sample is longer, we add 15 factors to our sample that were previously studied in the literature but not studied by Hou et al. (2020), and due to minor conservative factor

construction details that we believe robustify factor behavior.³ We discuss this decomposition further in Section 2, where we detail our factor construction choices and discuss why we prefer them.

The Hou et al. (2020) sample includes a number of factors that the original studies found to be insignificant.⁴ We exclude these when calculating the replication rate. After we make this adjustment, the replication rate rises to 61.3%, shown in the third bar in Figure 1.

Alpha, not raw return. Hou et al. (2020) analyze and test factors' raw returns, but if we wish to learn about "anomalies," economic theory dictates the use of risk-adjusted returns. Raw return gives a misleading inference for the factor if it differs from the alpha: When the raw return is significant, but the alpha is not, it simply means that the factor is taking risk exposure and the risk premium is significant, which does not indicate anomalous returns of the factor. Likewise, when the raw return is insignificant, but the alpha is significant, then the factor's efficacy is masked by its risk exposure. An example of this is the low-beta anomaly, where theory predicts that the alpha of a dollar-neutral low-beta factor is positive, but its raw return is negative or close to zero (Frazzini and Pedersen, 2014). In this case, the "failure to replicate" of Hou et al. (2020) is, in fact, support for the betting-against-beta theory. We analyze alpha to the CAPM, which is the clearest theoretical benchmark model that is not mechanically linked to other so-called anomalies in the list of replicated factors. The fourth bar in Figure 1 shows that the replication rate rises to 82.4% based on tests of factors' CAPM alpha.

Multiple testing and our Bayesian model. The first four bars in Figure 1 are based on individual ordinary least squares (OLS) t -statistics for each factor. But Harvey et al. (2016) rightly point out that this type of analysis suffers from a multiple testing (MT) problem. Harvey et al. (2016) recommend MT adjustments that raise the threshold for a t -statistic to be considered statistically significant. We report one such MT correction using

³We use tercile spreads while they use deciles; we use tercile breakpoints from all stocks above the NYSE 20th percentile (i.e., non-micro-caps), they use straight NYSE breakpoints; we always lag accounting data four months, they use a mixture of updating schemes; we exclude IBES factor due to their relatively short history; we use capped value-weighting they use straight value-weights; We look at returns over a 1 month holding period, they use 1, 6 and 12 months. In appendix C we detail how each change affects the replication rate.

⁴We identify 34 factors from Hou et al. (2020) for which the original paper did not find a significant alpha or did not study factor returns (see appendix Table J.1).

a leading method proposed by [Benjamini and Yekutieli \(2001\)](#). Accounting for MT in this manner, we find that the replication rate drops to 75.6% (the fifth bar of [Figure 1](#)). For comparison, [Hou et al. \(2020\)](#) consider a similar adjustment and find that their replication rate drops from 35% with OLS to 18% after MT correction.

However, common frequentist MT corrections can be unnecessarily crude. Our handling of the MT problem is different. We propose a Bayesian framework for the joint behavior of all the factors, resulting in an MT correction that sacrifices much less power than its frequentist counterpart (which we demonstrate via simulation).⁵ To understand the benefits of our approach, note first that we impose a prior that all alphas are expected to be zero. The role of the Bayesian prior is conceptually similar to that of frequentist MT corrections—it imposes conservatism on statistical inference and controls the false discovery rate. Second, our *joint* factor model allows us to conduct inference for all factor alphas simultaneously. The joint structure among factors leverages dependence in the data in order to draw more informative statistical inferences (relative to conducting independent individual tests). Our zero-alpha prior shrinks alpha estimates of all factors, thereby leading to fewer discoveries (i.e., a lower replication rate), with similar conservatism as a frequentist MT correction. At the same, however, the model allows us to learn more about the alpha of any individual factor, borrowing estimation strength across all factors. The improved precision of alpha estimates for all factors can increase the number of discoveries. Which effect dominates when we construct our final Bayesian model—the conservative shrinkage to the prior or the improved precision of alphas—is an empirical question.

In our sample, we find that the two effects exactly offset, which is why the Bayesian multiple testing view delivers a replication rate identical to the OLS-based rate. Specifically, our estimated replication rate rises to 82.4% (the sixth bar of [Figure 1](#)) using our Bayesian approach to the MT problem.⁶ The intuition behind this surprising result is simply that

⁵A large statistics literature (see [Gelman et al., 2013](#), and references therein) explains how Bayesian estimation naturally combats MT problems and [Gelman et al. \(2012\)](#) conclude that “the problem of multiple comparisons can disappear entirely when viewed from a hierarchical Bayesian perspective.” [Chinco et al. \(2020\)](#) use a Bayesian estimation framework similar to ours for a different (but conceptually related) problem. They infer the distribution of coefficients in a stock return prediction model to calculate what they dub the “anomaly base rate.”

⁶Our Bayesian approach leads to an even larger increase in the replication rate when using pure value-weighted returns (see [Figure C.1](#) of the appendix) and when considering global evidence outside the US (as

having many factors (a “factor zoo”) can be a strength rather than a weakness when assessing the replicability of factor research. It is obvious that our posterior is tighter when a factor has performed better and has a longer time series. But the posterior is further tightened if similar factors have also performed well, and if additional data shows that these factors have performed well in many other countries.⁷

Benefits of our model beyond the replication rate. One of the key benefits of Bayesian statistics is that one recovers not just a point estimate but the entire posterior distribution of parameters. The posterior allows us to make any possible probability calculation about parameters. For example, in addition to the replication rate, we also calculate the posterior probability of false discoveries (false discovery rate, FDR) and the posterior expected fraction of true factors. Moreover, we calculate Bayesian confidence intervals (also called credibility intervals) for each of these estimates. We find that our 82.4% replication rate has a tight posterior standard error of 2.8%. The posterior Bayesian FDR is only 0.1% with a 95% confidence interval of [0.0%, 1.0%], demonstrating the small risk of false discoveries. The expected fraction of true factors is 94.0% with a posterior standard error of 1.3%.

Global replication. Having found a high degree of internal validity of prior research, we next consider external validity across countries and over time. Regarding the former, we investigate how our conclusions are affected when we extend the data to include all factors in a large global panel of 93 countries. The last bar in Figure 1, shows that, based on the global sample, the final replication rate is 82.4%. This estimate is based on the Bayesian model extended to incorporate the joint behavior of international data. Because it accounts for the global correlation structure among factors, the model recognizes that international evidence is not independent out-of-sample evidence, and uses only the incremental global evidence to update the overall replicability assessment. And it continues to account for multiple testing. The global result reflects that factor performance in the US replicates well in an extensive cross section of countries. Serving as our final estimate, the global factor replication rate

we show later, in Figure 6).

⁷Taking this intuition further, we can glean additional information from studying whether factors work in other asset classes, as has been done for value and momentum (Asness et al., 2013), betting against beta (Frazzini and Pedersen, 2014), time series momentum (Moskowitz et al., 2012), and carry (Kojien et al., 2018).

more than doubles that of [Hou et al. \(2020\)](#) by grounding our tests in economic theory and modern Bayesian statistics. We conclude from the global analysis that factor research demonstrates external validity in the cross section of countries.

Post-publication performance. [McLean and Pontiff \(2016\)](#) find that US factor returns “are 26% lower out-of-sample and 58% lower post-publication.”⁸ Our Bayesian framework shows that, given a prior belief of zero alpha but an OLS alpha ($\hat{\alpha}$) that is positive, then our posterior belief about alpha lies somewhere between zero and $\hat{\alpha}$. Hence, a positive but attenuated post-publication alpha is the expected outcome based on Bayesian learning, rather than a sign of non-reproducibility. Further, when comparing factors cross-sectionally, the prediction of the Bayesian framework is that higher pre-publication alphas, if real, should be associated with higher post-publication alphas on average. And that is what we find. We present new and significant cross-sectional evidence that factors with higher in-sample alpha generally have higher out-of-sample alpha. The attenuation in the data is somewhat stronger than predicted by our Bayesian model. We conclude that factor research demonstrates external validity in the time series, but there appears to be some decay of the strongest factors that could be due to arbitrage or data mining.⁹

Publication bias. We also address the issue that factors with strong in-sample performance are more likely to be published while poorly performing factors are more likely to be unobserved in the literature. Publication bias can influence our full-sample Bayesian evidence through the empirical Bayes estimation of prior hyperparameters. To account for this bias, we show how to pick a prior distribution that is unaffected by publication bias by using only out-of-sample data or estimates from [Harvey et al. \(2016\)](#). Using such priors, the full-sample alphas are shrunk more heavily toward zero. The result is a slight drop in US the replication rate to 81.5%. If we add an extra degree of conservatism to the prior, the replication rate drops to 79.8%. Further, our out-of-sample evidence across time and across countries is not subject to publication bias.

⁸Extending the evidence to global stock markets, [Jacobs and Müller \(2020\)](#) find that “the United States is the only country with a reliable post-publication decline in long-short returns.” [Chen and Zimmermann \(2020b\)](#) use Bayesian methods to estimate bias-corrected post-publication performance and find that average returns drop by only 12% after publication in US data.

⁹Data prior to the sample used in original studies also constitutes out-of-sample evidence ([Linnainmaa and Roberts, 2018](#); [Ilmanen et al., 2021](#)). Our external validity conclusions are the same when we also include pre-original-study out-of-sample evidence.

Multidimensional challenge: A Darwinian view of the factor zoo. [Harvey et al. \(2016\)](#) challenge the sheer number of factors and [Cochrane \(2011\)](#) refers to as “the multidimensional challenge.” We argue that the factor research universe should not be viewed as hundreds of distinct factors. Instead, factors cluster into a relatively small number of highly correlated themes, and this property features prominently in our Bayesian modeling approach. We propose a factor taxonomy that algorithmically classifies factors into 13 themes possessing a high degree of within-theme return correlation and economic concept similarity, and low across-theme correlation. The emergence of themes, in which factors are minor variations on a related idea, is intuitive. For example, each value factor is defined by a specific valuation ratio, but there are many plausible ratios. Considering their variations is not spurious alpha-hacking, particularly when the “correct” value signal construction is debatable.

We estimate a replication rate of greater than 50% in 11 of the 13 themes (based on the Bayesian model including MT adjustment), the exceptions being “low leverage,” and “size” factor themes. We also analyze which themes matter when simultaneously controlling for all other themes. To do so, we estimate the ex post tangency portfolio of 13 theme-representative portfolios. We find that 10 of the 13 themes enter into the tangency portfolio with significantly positive weights, where the three displaced themes are “profitability,” “investment,” and “size.”

Why, the profession asks, have we arrived at a “factor zoo”?¹⁰ The answer, evidently, is because the risk-return tradeoff is complex and difficult to measure. The complexity manifests in our inability to isolate a single, silver bullet characteristic that pins down the risk-return tradeoff. Classifying factors into themes, we trace the economic culprits to roughly a dozen concepts. This is already a multidimensional challenge, but it is compounded by the fact that within a theme there are many detailed choices for how to configure the economic concept, which results in highly correlated within-theme factors. Together, the themes (and the factors in them) each make slightly different contributions to our collective understanding of markets. A more positive take on the factor zoo is *not* as a collective exercise in data

¹⁰See [Bryzgalova et al. \(2019\)](#), [Chordia et al. \(2020\)](#), [Kelly et al. \(2019\)](#), [Kozak et al. \(2020\)](#), [Green et al. \(2017\)](#), and [Feng et al. \(2020\)](#) for other perspectives on high-dimensional asset pricing problems, and [Chen \(2020\)](#) for an argument why *p*-hacking cannot explain the existence of so many significant factors.

mining and false discovery; instead, it is a natural outcome of a decentralized effort in which researchers make contributions that are correlated with, but incrementally improve on, the shared body of knowledge.

Economic implications. Our findings have broad implications for finance researchers and practitioners. We confirm that the body of finance research contains a multitude of replicable information about the drivers of expected returns. Further, we show that investors would have profited from factors deemed significant by our Bayesian method, but deemed insignificant by the frequentist MT method proposed by [Harvey et al. \(2016\)](#). Indeed, [Figure 2](#) plots the returns of the subset of factors discovered by our method but discarded by the frequentist method. These factors produce an annualized information ratio (IR) of 0.93 in the US and 1.10 globally (ex. US) over the full sample, with t -statistics above five. If we restrict analysis to the sample after that of [Harvey et al. \(2016\)](#), the performance differential remains large and significant. These findings show strong external validity (post original publications, post [Harvey et al. \(2016\)](#), different countries) and significant economic benefits of exploiting the joint information in all factor returns rather than simply applying a high cutoff for t -statistics.¹¹ We also show how the optimal risk-return profile has improved over time as factors have been discovered. In other words, the Sharpe ratio of the tangency portfolio has meaningfully increased over time as truly novel drivers of returns have been discovered. These findings can help inform asset pricing theory.

1 A Bayesian Model of Factor Replication

This section presents our Bayesian model for assessing factor replicability. We first draw out some basic implications of the Bayesian framework for interpreting evidence on individual factor alphas, then present a hierarchical structure for simultaneously modeling factors in a variety of themes and across many countries.

¹¹The out-of-sample performance across all significant factors under empirical Bayes is also highly significant as shown in [Appendix Figure D.1](#).




Figure 2: Out-of-sample performance of marginally significant factors

Note: The figure shows the cumulative CAPM alpha of an average of factors significant under our empirical Bayes framework, but not with the Benjamini-Yekutieli adjustment suggested by [Harvey et al. \(2016\)](#). The significance cutoffs are re-estimated each year with the available data. Factors are eligible for inclusion after the sample period in the original paper, so all returns are out-of-sample. The table shows the information ratio (alpha divided by residual volatility) for the full sample (1990-2020) and the post-[Harvey et al. \(2016\)](#) sample (2013-2020) with t -statistics in parentheses. The dashed line is at December 2012.

1.1 Learning About Alpha: The Bayes Case

Posterior Alpha

We begin by considering an excess return factor f_t . A study of “anomalous” factor returns requires a risk benchmark, without which we cannot separate distinctive factor behavior from run of the mill risk compensation. We assume a CAPM benchmark due to its history as a factor research benchmark for decades, and because it is not mechanically related to any of the factors that we attempt to replicate (in contrast to, say, the model of [Fama and French, 1993](#), which by construction explains size and value factors). A factor’s net performance versus the excess market factor (r_t^m) is its α :

$$f_t = \alpha + \beta r_t^m + \varepsilon_t. \tag{1}$$

Our Bayesian prior is that the alpha is normally distributed with mean zero and variance τ^2 , or $\alpha \sim N(0, \tau^2)$. The mean of zero implies that CAPM holds on average, and τ governs potential deviations from CAPM. Intuitively, the higher the confidence in the prior, the lower is τ . The error term, $\varepsilon_t \sim N(0, \sigma^2)$, has volatility σ , is independent and identically

distributed over time, and σ and β are observable.¹²

The risk-adjusted return, α , is estimated as the average market-adjusted factor return from T periods of data:

$$\hat{\alpha} = \frac{1}{T} \sum_t (f_t - \beta r_t^m) = \alpha + \frac{1}{T} \sum_t \varepsilon_t. \quad (2)$$

This observed ordinary least squares (OLS) estimate $\hat{\alpha}$ is distributed $N(\alpha, \sigma^2/T)$ given the true alpha, α . From Bayes' rule, we can compute the posterior distribution of the true alpha given the data evidence and prior. The posterior exhaustively describes the Bayesian's beliefs about alpha at a future time $t > T$ given the past experience, including the posterior expected factor performance,

$$E(\alpha|\hat{\alpha}) = E\left(f_t - \beta r_t^m \middle| \hat{\alpha}\right). \quad (3)$$

We derive the posterior alpha distribution via Bayes' rule (the derivation, which is standard, is shown in Appendix A). The posterior alpha is normal with mean

$$E(\alpha|\hat{\alpha}) = \kappa \hat{\alpha} \quad (4)$$

where κ is a shrinkage factor given by

$$\kappa = \frac{\tau^2}{\tau^2 + \sigma^2/T} = \frac{1}{1 + \frac{\sigma^2}{\tau^2 T}} \in (0, 1) \quad (5)$$

and the posterior variance is

$$\text{Var}(\alpha|\hat{\alpha}) = \kappa \frac{\sigma^2}{T} = \frac{1}{\frac{1}{\sigma^2/T} + \frac{1}{\tau^2}}. \quad (6)$$

The first insight from this posterior is that a Bayesian predicts future returns will have

¹²Here we seek to derive some simple expressions that illustrate the economic implications of Bayesian logic. In the empirical implementation, we use slightly richer model, as discussed further below. The empirical implementation normalizes factors so that σ is given at 10% for all factors, while β must be estimated, but this does not affect the economic points that we make in this section.

smaller alpha (in absolute value) than the OLS estimate $\hat{\alpha}$, because the posterior mean ($\kappa\hat{\alpha}$) must lie between $\hat{\alpha}$ and the prior mean of zero. Said differently, a large observed alpha might be due to luck and, given the prior, we expect that at least part of this performance indeed is luck. The more data we have (higher T), the less shrinkage there is (i.e., κ closer to 1). Likewise, the stronger is the prior of zero alpha (i.e., lower τ), the heavier is the shrinkage. We can think of the prior τ in terms of the number of time periods of evidence that it corresponds to. That is, the posterior mean, $E(\alpha|\hat{\alpha})$, corresponds to first observing σ^2/τ^2 time periods with an average alpha of zero, followed by T time periods with a average alpha of $\hat{\alpha}$.

When evaluating out-of-sample evidence, a positive, but lower, alpha is sometimes interpreted as a sign of replication failure. But this is the expected outcome from the Bayesian perspective (i.e., based on the latest posterior), and can be fully consistent with a high degree of replicability. In fact, we show later that the comparatively low post-publication factor performance documented by [McLean and Pontiff \(2016\)](#) turns out to be consistent with the posterior a Bayesian would have formed given published results. Thus, post-publication results have tended to confirm the Bayesian’s beliefs and as a result the Bayesian posterior alpha estimate has been extraordinarily stable over time (see Section 3.2).

Alpha-hacking

Because out-of-sample alpha attenuation is not generally a sign of replication failure, we may want a more direct probe for non-replicability. We can build such a test into our Bayesian framework by embedding scope for “alpha-hacking,” or selectively reporting or manipulating data to artificially make the alpha seem larger. We represent this idea using the following distribution of factor returns in the in-sample time period $t = 1, \dots, T$:

$$f_t = \alpha + \beta r_t^m + \underbrace{\tilde{\varepsilon}_t + u}_{\varepsilon_t}. \tag{7}$$

Here, $\tilde{\varepsilon}_t \sim N(0, \sigma^2)$ captures usual return shocks and $u \sim N(\bar{\varepsilon}, \sigma_u^2)$ represents return inflation due to alpha-hacking. The total in-sample return shock ε_t is normally distributed, $N(\bar{\varepsilon}, \bar{\sigma}^2)$, where $\bar{\varepsilon} \geq 0$ is the alpha-hacking bias, and the variance $\bar{\sigma}^2 = \sigma^2 + \sigma_u^2 \geq \sigma^2$ is elevated due to

the artificial noise created by alpha-hacking.¹³ Naturally, the false benefits of alpha-hacking disappear in out-of-sample data, or in other words $\varepsilon_t \sim N(0, \sigma^2)$ for $t > T$. The Bayesian accounts for alpha-hacking as follows:

Proposition 1 (Alpha-hacking) *The posterior alpha with alpha-hacking is given by*

$$E(\alpha|\hat{\alpha}) = -\kappa_0 + \kappa^{\text{hacking}} \hat{\alpha} \tag{8}$$

where $\kappa^{\text{hacking}} = \frac{1}{1 + \frac{\sigma^2}{\tau^2 T}} \leq \kappa$ and $\kappa_0 = \kappa^{\text{hacking}} \bar{\varepsilon} \geq 0$. Further, $\kappa^{\text{hacking}} \rightarrow 0$ in the limit of “pure alpha-hacking,” $\tau \rightarrow 0$ or $\bar{\sigma} \rightarrow \infty$.

The Bayesian posterior alpha accounts for alpha-hacking in two ways. First, the estimated alpha is shrunk more heavily toward zero since the factor κ^{hacking} is now smaller. Second, the alpha is further discounted by the intercept term κ_0 due to the bias in the error terms.

We examine alpha-hacking empirically in Section 3.2 in light of Proposition 1. We consider a cross-sectional regression of factors’ out-of-sample (e.g., post-publication) alphas on their in-sample alphas, looking for the signatures of alpha-hacking in the form of a negative intercept term or a slope coefficient that is too small. In addition, Section 3.3 shows how to estimate the Bayesian model in a way that is less susceptible to the effects of alpha hacking and Appendix A presents additional theoretical results characterizing alpha-hacking.

1.2 Hierarchical Bayesian Model

Shared Alphas: The Case of Complete Pooling

We now embed a critical aspect of factor research into our Bayesian framework: Factors are often correlated and conceptually related to each other. For concreteness, we begin with a setting in which the researcher has access to “domestic” evidence in (1) as well as “global” evidence from an international factor, f_t^g , with known exposure β^g to the global market index r_t^g :

$$f_t^g = \alpha + \beta^g r_t^g + \varepsilon_t^g. \tag{9}$$

¹³We note that this elevated variance cannot be detected by looking at the in-sample variance of residual returns since the alpha-hacking term u does not depend on time t .

Here, we assume that the true alpha for this global factor is the same as the domestic alpha. In other words, we have complete “pooling” of information about alpha across the two samples. As an alternative interpretation, the researcher could have access to two related factors, say two different value factors in the same country, and assume that they have the same alpha because they capture the same investment principle.

The global shock, ε_t^g , is normally distributed $N(0, \sigma^2)$, and ε_t^g and ε_t are jointly normal with correlation ρ .¹⁴ The estimated alpha based on the global evidence is simply its market-adjusted return:

$$\hat{\alpha}^g = \frac{1}{T} \sum_t (f_t^g - \beta^g r_t^g). \quad (10)$$

To see the power of global evidence (or, more generally, the power of observing related strategies), we consider the posterior when observing both the domestic and global evidence.

Proposition 2 (The Power of Shared Evidence) *The posterior alpha given the domestic estimate, $\hat{\alpha}$, and the global estimate, $\hat{\alpha}^g$, is normally distributed with mean*

$$E(\alpha | \hat{\alpha}, \hat{\alpha}^g) = \kappa^g \left(\frac{1}{2} \hat{\alpha} + \frac{1}{2} \hat{\alpha}^g \right). \quad (11)$$

The global shrinkage parameter is

$$\kappa^g = \frac{1}{1 + \frac{\sigma^2}{\tau^2 T} \frac{1+\rho}{2}} \in [\kappa, 1] \quad (12)$$

which decreases with the correlation ρ , attaining the minimum value, $\kappa^g = \kappa$, when $\rho = 1$.

The posterior variance is lower when observing both domestic and global evidence:

$$\text{Var}(\alpha | \hat{\alpha}) \geq \text{Var}(\alpha | \hat{\alpha}, \hat{\alpha}^g). \quad (13)$$

Naturally, the posterior depends on the average alpha observed domestically and globally. Furthermore, the combined alpha is shrunk toward the prior of zero. The shrinkage factor κ^g

¹⁴The framework can be generalized to a situation where the global shocks have a different volatility and sample length. In this case, the Bayesian posterior puts more weight on the sample with lower volatility and longer sample.

is smaller (heavier shrinkage) if the markets are more correlated because the global evidence provides less new information. With low correlation, the global evidence adds a lot of independent information, shrinkage is lighter, and the Bayesian becomes more confident in the data and less reliant on the prior. The proposition shows that, if a factor has been found to work both domestically and globally, then the Bayesian expects stronger out-of-sample performance than a factor that has only worked domestically (or has only been analyzed domestically).

Two important effects are at play here, and both are important for understanding the empirical evidence presented below: The domestic and global alphas are shrunk both toward *each other* and toward *zero*. For example, suppose that a factor worked domestically but not globally, say $\hat{\alpha} = 10\% > \hat{\alpha}^g = 0\%$. Then the overall evidence points to an alpha of $\frac{1}{2}\hat{\alpha} + \frac{1}{2}\hat{\alpha}^g = 5\%$, but shrinkage toward the prior results in a lower posterior, say, 2.5%. Hence, the Bayesian expects future factor returns in both regions of 2.5%. That shared alphas are shrunk together is a key feature of a *joint* model, and it generally leads to different conclusions than when factors are evaluated independently. Next we consider a perhaps more realistic model in which factors are only partially shrunk toward each other.

Hierarchical Alphas: The Case of Partial Pooling

We now consider several factors, numbered $i = 1, \dots, N$. Factor i has a true alpha given by

$$\alpha^i = c + w^i. \tag{14}$$

Here, c is the common component of all alphas, which has a prior distribution given by $N(0, \tau_c^2)$. Likewise, w^i is the idiosyncratic alpha component, which has a prior distribution given by $N(0, \tau_w^2)$, independent of c and across i . Said differently, we can imagine that nature first picks of the overall c from $N(0, \tau_c^2)$ and then picks the factor-specific α^i from $N(c, \tau_w^2)$.

This hierarchical model is a realistic compromise between assuming that all factor alphas are completely different (using equation (4) for each alpha separately) and assuming that they are all the same (using Proposition 2). Rather than assuming no pooling or complete pooling, the hierarchical model allows factors to have a common component and an idiosyncratic

component.

Suppose we observe factor returns of

$$f_t^i = \alpha^i + \beta^i r_t^m + \varepsilon_t^i \quad (15)$$

where ε_t^i are normally distributed with mean 0 and variance σ^2 and $\text{Cor}(\varepsilon_t^i, \varepsilon_t^j) = \rho \geq 0$ for all i, j .¹⁵ Computing the observed alpha estimates as above, $\hat{\alpha}^i = \frac{1}{T} \sum_t (f_t^i - \beta^i r_t^m)$, we derive the posterior in the following result.¹⁶

Proposition 3 (Hierarchical Alphas) *The posterior alpha of factor i given the evidence on all factors is normally distributed with mean*

$$E(\alpha^i | \hat{\alpha}^1, \dots, \hat{\alpha}^N) = \frac{1}{1 + \frac{\rho\sigma^2}{\tau_w^2 T} + \frac{\tau_w^2 + (1-\rho)\sigma^2/T}{\tau_c^2 N}} \hat{\alpha}^i + \frac{1}{1 + \frac{(1-\rho)\sigma^2}{\tau_w^2 T}} \left(\hat{\alpha}^i - \frac{1}{1 + \frac{\tau_w^2 + (1-\rho)\sigma^2/T}{(\tau_c^2 + \rho\sigma^2/T)N}} \hat{\alpha}^\cdot \right), \quad (16)$$

where $\hat{\alpha}^\cdot = \frac{1}{N} \sum_j \hat{\alpha}^j$ is average alpha. When the number of factors N grows, the limit is

$$\lim_{N \rightarrow \infty} E(\alpha^i | \hat{\alpha}^1, \dots, \hat{\alpha}^N) = \frac{1}{1 + \frac{\rho\sigma^2}{\tau_w^2 T}} \hat{\alpha}^i + \frac{1}{1 + \frac{(1-\rho)\sigma^2}{\tau_w^2 T}} (\hat{\alpha}^i - \hat{\alpha}^\cdot). \quad (17)$$

The posterior variance of factor i 's alpha using the information in all factor returns is lower than the posterior variance when looking at this factor in isolation:

$$\text{Var}(\alpha^i | \hat{\alpha}^1, \dots, \hat{\alpha}^N) < \text{Var}(\alpha^i | \hat{\alpha}^i). \quad (18)$$

¹⁵Alternatively, we can write the error terms in a similar way to how we write the alphas in (14), namely $\varepsilon_t^i = \sqrt{\rho} \tilde{\varepsilon}_t + \sqrt{1-\rho} \hat{\varepsilon}_t^i$, where $\tilde{\varepsilon}_t$ are idiosyncratic shocks that are independent across factors and of the common shock $\hat{\varepsilon}_t$, with $\text{Var}(\tilde{\varepsilon}_t) = \text{Var}(\hat{\varepsilon}_t) = \sigma^2$. We note that we require (the empirically realistic case) that $\rho \geq 0$ since we cannot have an arbitrarily large number of normal random variables with equal negative correlation (because the corresponding variance-covariance matrix would not be positive semi-definite for large enough N).

¹⁶The general hierarchical model is used extensively in the statistics literature, see, e.g., [Gelman et al. \(2013\)](#), but to our knowledge the results in Proposition 3 are not in the literature.

The posterior variance is decreasing in N and, as $N \rightarrow \infty$, its limit is

$$\text{Var}(\alpha^i | \hat{\alpha}^1, \dots, \hat{\alpha}^N) \searrow \frac{\rho\sigma^2}{T} \frac{1}{1 + \frac{\rho\sigma^2}{\tau_c^2 T}} + \frac{(1-\rho)\sigma^2}{T} \frac{1}{1 + \frac{(1-\rho)\sigma^2}{\tau_w^2 T}}. \quad (19)$$

The main insight of this proposition is that having data on many factors is helpful for estimating the alpha of any of them. Intuitively, the posterior for any individual alpha depends on all of the other observed alphas because they are all informative about the common alpha component. That is, the other observed alphas tell us whether alpha exists in general or, said another way, tell us if the CAPM appears to be violated in general. Further, the factor's own observed alpha tells us whether this specific factor appears to be especially good or bad. Using all of the factors jointly reduces posterior variance for all alphas. In summary, the joint model with hierarchical alphas has the dual benefits of identifying the common component in alphas and tightening confidence intervals by sharing information among factors.

To understand the proposition in more detail, consider first the (unrealistic) case in which all factor returns have independent shocks ($\rho = 0$). In this case, we essentially know the overall alpha when we see many uncorrelated factors. Indeed, the average observed alpha becomes a precise estimator of the overall alpha with more and more observed factors, $\hat{\alpha} \rightarrow c$. Since we essentially know the overall alpha in this limit, the first term in (17) becomes $1 \times \hat{\alpha}$ when $\rho = 0$ meaning that we don't need any shrinkage here. The second term is the outperformance of factor i above the average alpha, and this outperformance is shrunk toward our prior of zero. Indeed, the outperformance is multiplied by a number less than one, and this multiplier naturally decreases in the return volatility σ and decreases in our conviction in the prior (increases in τ_w).

The posterior variance is also intuitive in the case of $\rho = 0$. The posterior variance is clearly lower compared to only observing the performance of factor i itself:

$$\text{Var}(\alpha^i | \hat{\alpha}^1, \hat{\alpha}^2, \dots) = \frac{\sigma^2}{T} \frac{1}{1 + \frac{\sigma^2}{\tau_w^2 T}} < \frac{\sigma^2}{T} \frac{1}{1 + \frac{\sigma^2}{(\tau_c^2 + \tau_w^2) T}} = \text{Var}(\alpha^i | \hat{\alpha}^i) \quad (20)$$

based on (19) and (6). With partial pooling, the posterior variance decreases because the

denominator on the left does not have τ_c^2 , reflecting that uncertainty about the general alpha has been eliminated by observing many factors.

In the realistic case where factor returns are correlated ($\rho > 0$), we see that both the average alpha $\hat{\alpha}$ and factor i 's outperformance $\hat{\alpha}^i - \hat{\alpha}$ are shrunk toward the prior of zero. This is because we cannot precisely estimate the overall alpha even with an infinite number of correlated factors—the correlated part never vanishes. Nevertheless, we still shrink the confidence interval, $\text{Var}(\alpha^i | \hat{\alpha}^1, \dots, \hat{\alpha}^N) \leq \text{Var}(\alpha^i | \hat{\alpha}^i)$, since more information is always better than less.

Multi-level Hierarchical Model

The model development to this point is simplified to draw out its intuition. Our empirical implementation is based on a more realistic (and slightly more complex) model that takes into account that factors naturally belong to different economic themes and to different regions.

In our global analysis, we have N different characteristic signals (e.g., book-to-market) across K regions, for a total of NK factors (e.g., US, developed, and emerging markets versions of book-to-market). Each of the N signals belongs to a smaller number of J theme clusters, where one cluster consists of various value factors, another consists of various momentum factors, and so on. One level of our hierarchical model allows for partially shared alphas among factors in the same theme cluster. Another level allows for commonality across regions among factors associated with the same underlying characteristic, capturing for example the connection between the book-to-market factor in different markets.

Mathematically, this means that an individual factor i has an alpha of

$$\alpha^i = \alpha^o + c^j + s^n + w^i. \tag{21}$$

Concretely, suppose factor $i \in \{1, \dots, NK\}$ is the book-to-market factor in the US region. Part of its alpha is driven by a component that is common to all factors, α^o , which we dogmatically fix at zero to be conservative. In addition, this factor i belongs to the value cluster $j \in \{1, \dots, J\}$, which contributes a cluster-specific alpha $c^j \sim N(0, \tau_c^2)$. Next, since

factor i is based on book-to-market characteristic $n \in \{1, \dots, N\}$, it has an incremental signal-specific alpha of $s^n \sim N(0, \tau_s^2)$ that is shared across regions—e.g., it’s the common behavior among book-to-market factors regardless of geography. Finally, $w^i \sim N(0, \tau_w^2)$ is factor i ’s idiosyncratic alpha, namely the incremental alpha that is unique to the US version of book-to-market.

We write this model in vector form as¹⁷

$$\alpha = \alpha^o \mathbf{1}_{NK} + Mc + Zs + w \quad (22)$$

where $\alpha = (\alpha^1, \dots, \alpha^{NK})'$, $c = (c^1, \dots, c^J)'$, $s = (s^1, \dots, s^N)'$, $w = (w^1, \dots, w^{NK})'$, M is the $NK \times J$ matrix of cluster memberships, and Z is the $NK \times N$ matrix indicating the characteristic that factor i is based on. In particular, $M_{i,j} = 1$ if factor i is in cluster j and $M_{i,j} = 0$ otherwise. Likewise, $Z_{i,n} = 1$ if factor i is based on characteristic n and $Z_{i,n} = 0$ otherwise. This hierarchical model implies that the prior variance of alpha, denoted Ω , is¹⁸

$$\Omega = \text{Var}(\alpha) = MM'\tau_c^2 + ZZ'\tau_s^2 + I_{NK}\tau_w^2. \quad (23)$$

In some cases, we analyze this model within a single region, $K = 1$ (for example, in our US-only analysis). In this case, there is no difference between signal-specific alphas and idiosyncratic alphas, so we collapse one level of the model by setting $\tau_s = 0$ and $s^n = 0$ for $n \in \{1, \dots, N\}$. In any case, the following result shows how to compute the posterior distribution of all alphas based on the prior uncertainty, Ω , and a general variance-covariance matrix of return shocks, $\Sigma = \text{Var}(\varepsilon)$. This result is at the heart of our empirical analysis.

Proposition 4 *In the multi-level hierarchical model, the posterior of the vector of true alphas is normally distributed with posterior mean*

$$E(\alpha|\hat{\alpha}) = (\Omega^{-1} + T\Sigma^{-1})^{-1} (\Omega^{-1}\mathbf{1}_{NK}\alpha_0 + T\Sigma^{-1}\hat{\alpha}) \quad (24)$$

¹⁷The notation $\mathbf{1}_N$ refers to an $N \times 1$ vector of ones and I_N is the $N \times N$ identity matrix.

¹⁸Stated differently, each diagonal element of Ω is $\tau_c^2 + \tau_s^2 + \tau_w^2$. Further, if $i \neq k$, then the $(i, k)^{th}$ element of Ω is $\tau_c^2 + \tau_s^2$ if i and k are constructed from the same signal in the same cluster in different regions, it is τ_c^2 if i and k are constructed from different signals in the same cluster, and it is 0 if i and k are in different clusters.

and posterior variance

$$\text{Var}(\alpha|\hat{\alpha}) = (\Omega^{-1} + T\Sigma^{-1})^{-1}. \quad (25)$$

As noted above, we set the mean prior alpha to zero ($\alpha_0 = 0$) in our empirical implementation. This prior is based on economic theory and leads to a conservative shrinkage toward zero as seen in (24). We note that, in the data, the observed alphas are mostly positive, not centered around zero. However, these positive alphas are related to the way that factors are signed, namely according to the convention in the original paper, which almost always leads to a positive factor return in the original sample. However, if we view this signing convention as somewhat arbitrary, then a symmetry argument implies that a prior of zero is again natural. Said differently, factor means would be centered around zero if we changed signs arbitrarily, so our prior is agnostic about these signs.

1.3 Bayesian Multiple Testing and Empirical Bayes Estimation

Frequentist MT corrections embody a principle of conservatism that seeks to limit false discoveries, controlling the family-wise error rate (FWER) or the false discovery rate (FDR). Leading frequentist methods achieve this by widening confidence intervals and raising p -values, but do not alter the underlying point estimate.

Bayesian Multiple Testing

A large statistics literature describes how Bayesian modeling is effective for making reliable inferences in the face of multiple testing.¹⁹ Drawing on this literature, our hierarchical model is a prime example of how Bayesian methods accomplish their MT correction based on two key model features.

First is the model prior, which imposes statistical conservatism in analogy to frequentist MT methods. It anchors the researcher's beliefs to a sensible default (e.g., all alphas are zero) in case the data are insufficiently informative about the parameters of interest. Reduction

¹⁹See Gelman et al. (2012); Berry and Hochberg (1999); Greenland and Robins (1991); Efron and Tibshirani (2002), among others. See Gelman (2016) for an intuitive, informal discussion of the topic.

of false discoveries is achieved first by shrinking estimates toward the prior. When there is no information in the data, the alpha point estimate is the prior mean and there are no false discoveries. As data evidence accumulates, posterior beliefs migrate away from the prior toward the OLS alpha estimate. In the process, discoveries begin to emerge, though they remain dampened relative to OLS. In the large data limit, Bayesian beliefs converge on OLS with no MT correction, which is justified because in the limit there are no false discoveries. In other words, the prior embodies a particularly flexible form conservatism—the Bayesian model decides how severe of an MT correction to make based on the informativeness of the data.

Second is the hierarchical structure that captures joint behavior of factors. Modeling factors jointly means that each alpha is shrunk toward its cluster mean (i.e., toward related factors), in addition to being shrunk toward the prior of zero. So, if we observe a cluster of factors in which most perform poorly, then this evidence reduces the posterior alpha even for the few factors with strong performance—another form of Bayesian MT correction. In addition to this Bayesian discovery control coming through shrinkage of the posterior mean alpha, the Bayesian confidence interval also plays an important role and changes as a function of the data. Indeed, having data on related factors leads to a contraction of the confidence intervals in our joint Bayesian model. So while alpha shrinkage often has the effect of reducing discoveries, the increased precision from joint estimation has the opposite effect of enhancing statistical power and thus increases discoveries.

In summary, a typical implementation of frequentist MT corrections estimates parameters independently for each factor, leaves these parameters unchanged, but inflates p -values to reduce the number of discoveries. In contrast, our hierarchical model leverages dependence in the data to efficiently learn about all alphas simultaneously. All data therefore helps to determine the center and width of each alpha’s confidence interval (Propositions 3 and 4). This leads to more precise estimates with “built-in” Bayesian MT correction.

Empirical Bayes Estimation

Given the central role of the prior, it might seem problematic that the severity of the Bayesian MT adjustment is at the discretion of the researcher. A powerful (and somewhat surprising)

aspect of a hierarchical model is that the prior can be learned in part from the data. This idea is formalized in the idea of “empirical Bayes (EB)” estimation, which has emerged as a major toolkit for navigating multiple tests in high-dimensional statistical settings (Efron, 2012).

The general approach to EB is to specify a multi-level hierarchical model, and then to use the dispersion of estimated effects within each level to learn about the prior parameters for that level. In our setting, the specific implementation of EB is dictated by Proposition 4. We first compute each factor’s abnormal return, $\hat{\alpha}$, as the intercept in a CAPM regression on the market excess return. Next, we set the overall alpha prior mean, α^o , to zero to enforce conservatism in our inferences.

From here, the benefits of EB kick in: The realized dispersion in alphas across factors helps to determine the appropriate level of conviction for the prior (that is, the appropriate values for τ_c^2 , τ_s^2 , and τ_w^2). For example, if we compute the average alpha for each cluster, \hat{c}^j (e.g., the average value alpha, the average momentum alpha, and so on), the cross-sectional variation in \hat{c}^j suggests that $\tau_c^2 \cong \frac{1}{J-1} \sum_{j=1}^J (\hat{c}^j - \hat{c})^2$. The same idea applies to τ_s^2 . Likewise, variation in observed alphas after accounting for hierarchical connections is informative about $\tau_w^2 \cong \frac{1}{NK-N-J} \sum_{i=1}^N (\hat{w}^i)^2$, where $\hat{w} = \hat{\alpha} - M\hat{c} - Z\hat{s}$.

The above variances illustrate the point that EB can help calibrate prior variances using the data itself. But those calculations are too crude, because they ignore sampling variation coming from the noise in returns, ε , which has covariance matrix Σ . Empirical Bayes estimates the prior variances by maximizing the prior likelihood function of the observed alphas, $\hat{\alpha} \sim N(0, \Omega(\tau_c, \tau_s, \tau_w) + \hat{\Sigma}/T)$, where the notation emphasizes that Ω depends on τ_c , τ_s , and τ_w according to (23). The likelihood function accounts for sampling variation through the a plug-in estimate of the covariance matrix of factor return shocks, $\hat{\Sigma}$.²⁰ We collect the resulting hyper-parameters in τ , that is, $\tau_c, \tau_s, \tau_w, \hat{\Sigma}$, and β^i .

Bayesian FDR and FWER

With the EB estimates (τ) on hand, we can compute the posterior distribution of the alphas from Proposition 4. From the posterior, we can in turn compute Bayesian versions of the

²⁰We discuss the details of our EB estimation procedure in Appendix B.

FDR and FWER. Suppose that we consider a factor to be “discovered” if its z -score is greater than the critical value $\bar{z} = 1.96$:

$$\frac{E(\alpha^i | \hat{\alpha}^1, \dots, \hat{\alpha}^N, \tau)}{\sqrt{\text{Var}(\alpha^i | \hat{\alpha}^1, \dots, \hat{\alpha}^N, \tau)}} > \bar{z}. \quad (26)$$

Equivalently, factor i is discovered if $p\text{-null}_i < 2.5\%$,²¹ where we use the posterior to compute

$$p\text{-null}_i = Pr(\alpha^i < 0 | \hat{\alpha}^1, \dots, \hat{\alpha}^N, \tau). \quad (27)$$

In words, $p\text{-null}_i$ is the posterior probability that the null hypothesis is true, which is the Bayesian version of a frequentist p -value. Said differently, it is the posterior probability of a “false discovery,” namely the probability that the true alpha is actually non-positive.

We can further compute the Bayesian FDR as:

$$\text{FDR}^{\text{Bayes}} = E \left(\frac{\sum_i 1_{\{i \text{ false discovery}\}}}{\sum_i 1_{\{i \text{ discovery}\}}} \middle| \hat{\alpha}^1, \dots, \hat{\alpha}^N, \tau \right) \quad (28)$$

where we condition on the data including at least one discovery (so the denominator is not zero), otherwise FDR is set to zero (see [Benjamini and Hochberg, 1995](#)).

The following proposition is a novel characterization of the Bayesian FDR, and shows that it is the posterior probability of a false discovery, averaged across all discoveries:

Proposition 5 (Bayesian FDR) *Conditional on the parameters of the prior distribution τ and data with at least one discovery, the Bayesian false discovery rate can be computed as:*

$$\text{FDR}^{\text{Bayes}} = \frac{1}{\#\text{discoveries}} \sum_{i \text{ discovery}} p\text{-null}_i. \quad (29)$$

and is bounded, $\text{FDR}^{\text{Bayes}} \leq 2.5\%$.

This result shows explicitly how the Bayesian framework controls the false discovery rate without the need for additional MT adjustments.²² The definition of a discovery ensures

²¹We use a critical value of 2.5% rather than 5% because the 1.96 cut-off corresponds to a 2-sided test, while false discoveries are only on one side in the Bayesian framework.

²²[Efron \(2007\)](#) includes related analysis but, to our knowledge, this particular result is new.

that at most 2.5% of the discoveries are false according to the Bayesian posterior, which is exactly the right distribution for assessing discoveries from the perspective of the Bayesian. Further, if many of the discovered factors are highly significant (as is the case in our data), then the Bayesian FDR is much lower than 2.5%.²³

We can also compute a Bayesian version of the family-wise error rate, which is the probability of making one or more false discoveries in total:

$$\text{FWER}^{\text{Bayes}} = Pr \left(\sum_i 1_{\{i \text{ false discovery}\}} \geq 1 \mid \hat{\alpha}^1, \dots, \hat{\alpha}^N, \tau \right). \quad (30)$$

If we define a discovery as in (26) using the standard critical value $\bar{z} = 1.96$, then we do not necessarily control the family-wise error rate, $\text{FWER}^{\text{Bayes}}$, which is a harsh criterion that is concerned with the risk of a single false discovery without regard for the number of missed discoveries. $\text{FWER}^{\text{Bayes}}$ is a probability that can be computed from the posterior so it is straightforward to choose a critical value \bar{z} to ensure $\text{FWER}^{\text{Bayes}} \leq 5\%$ or any other level one prefers. The main point is that the Bayesian approach to replication lends itself to any inferential calculation the researcher desires because the posterior is a complete characterization of Bayesian beliefs about model parameters.

A Comparison of Frequentist and Bayesian False Discovery Control

We illustrate the benefits of Bayesian inference for our replication analysis via simulation. We assume a factor generating process based on the hierarchical model above and, for simplicity, consider a single region (as in our empirical US-only analysis), removing s^n and τ_s^2 from equations (21) and (23). We analyze discoveries as we vary the prior variances τ_c and τ_w . The remaining parameters are calibrated to our estimates for the US region in our empirical analysis below.

We simulate an economy with 130 factors in 13 different clusters of 10 factors each, observed monthly over 70 years. We assume that the mean alpha, α^o , is zero. We then draw

²³Proposition 5 formalizes the argument of Greenland and Robins (1991) that “from the empirical-Bayes or Bayesian perspective, multiple comparisons are not really a ‘problem.’ Rather, the multiplicity of comparisons provides an opportunity to improve our estimates through judicious use of any prior information (in the form of model assumptions) about the ensemble of parameters being estimated.”

a cluster alpha from $c^j \sim N(0, \tau_c^2)$ and a factor-specific alpha as $w^i \sim N(0, \tau_w^2)$. Based on these alphas, we generate realized returns by adding Gaussian noise.²⁴

We compute p -values separately using OLS with no adjustment or adjusting with the Benjamini-Yekutieli (BY) method. We also use EB to estimate the posterior alpha distribution, treating τ_c and τ_w as known in order to simplify simulations and focus on the Bayesian updating. For OLS and BY, a discovery occurs when the alpha estimate is positive and the two-sided p -value is below 5%. For EB, we consider it a discovery when the posterior probability that alpha is negative is less than 2.5%. For each pair of τ_c and τ_w , we draw 10,000 simulated samples, and report average discovery rates over all simulations.

Figure 3 reports alpha discoveries based on the OLS, BY, and EB approaches. For each method, we report the true FDR in the top panels (we know the truth since this is a simulation) and the “true discovery rate”²⁵ in the bottom panels.

When idiosyncratic variation in true alphas is small (left panels with $\tau_w = 0.01\%$) and the variation in cluster alphas is also small (values of τ_c near zero on the horizontal axis), alphas are very small and true discoveries are unlikely. In this case, the OLS false discovery rate can be as high as 25% as seen in the upper left panel. However, both BY and EB successfully correct this problem and lower the FDR. The lower left panel shows that the BY correction pays a high price for its correction in terms of statistical power when τ_c is larger. In contrast, EB exhibits much better power to detect true positives while maintaining a similar false discovery control as BY. In fact, when there are more discoveries to be made in the data (as τ_c increases), EB becomes even more likely to identify true positives than OLS. This is due to the joint nature of the Bayesian model, whose estimates are especially precise compared to OLS due to EB’s ability to learn more efficiently from dependent data. This illustrates a point of Greenland and Robins (1991) that “Unlike conventional multiple comparisons, empirical-Bayes and Bayes approaches will alter and can improve point estimates and can

²⁴The noise covariance matrix has a block structure calibrated to our data, with a correlation of 0.58 among factors in the same cluster and a correlation of 0.02 across clusters. The residual volatility for each factor is 10% per annum.

²⁵We define the true discovery rate to be the number of significantly positive alphas according to, respectively, OLS, BY, and EB divided by the number of truly positive alphas. Given our simulation structure, half of the alphas are expected to be positive in any simulation. Some of these will be small (i.e., economically insignificant) positives, so a testing procedure would require a high degree of statistical power to detect them. This is why the true discovery rate is below one even for high values of τ_c .

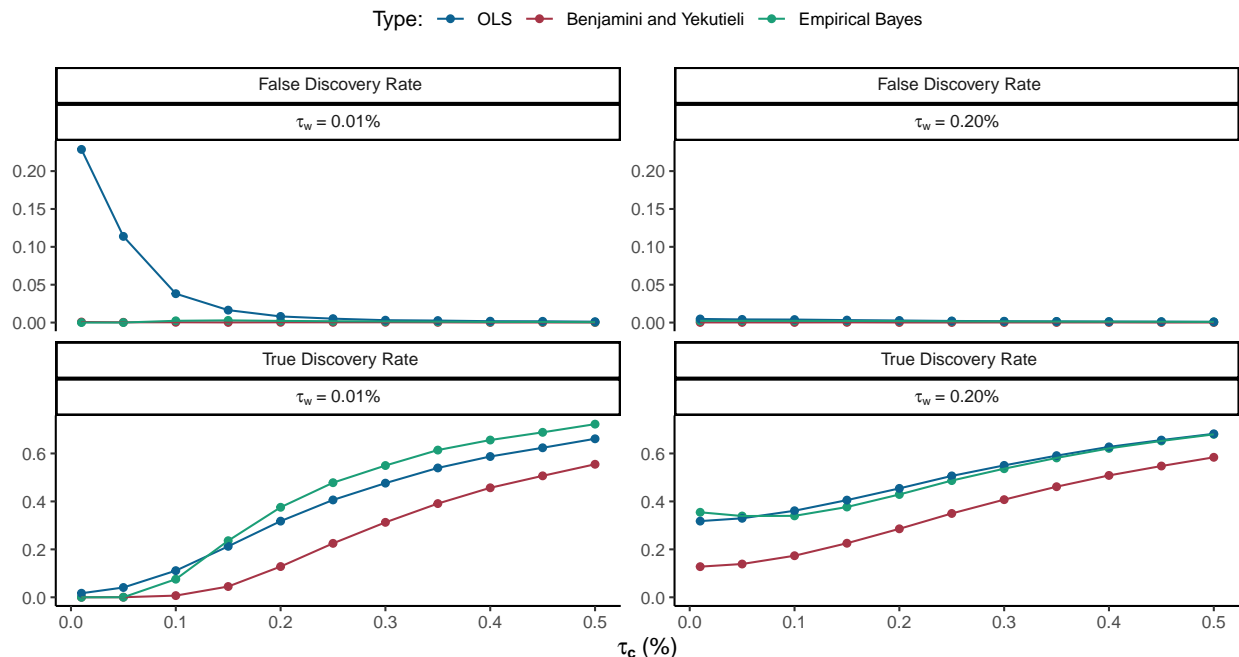


Figure 3: Simulation Comparison of False Discovery Rates

Note: The upper panels show the realized false discovery rate computed as the proportion of discovered factors for which the true alpha is negative, averaged over 10,000 simulations. The lower panels show the true discovery rate computed as the number of discoveries where the true alpha is positive divided by the total number of factors where the true alpha is positive. The left and right panels use low and high values of idiosyncratic variation in alphas (τ_w), respectively. The x -axis varies cluster alpha dispersion, τ_c .

provide more powerful tests and more precise (narrower) interval estimators.” When the idiosyncratic variation is larger ($\tau_w = 0.20\%$), there are many more true discoveries to be made, so the false discovery rate tends to be low even for OLS with no correction. Yet in the lower right panel we continue to see the costly loss of statistical power suffered by the BY correction.

In summary, EB accomplishes a flexible MT adjustment by adapting to the data generating process. When discoveries are rare so that there is a comparatively high likelihood of false discovery, EB imposes heavy shrinkage and behaves similarly to the conservative BY correction. In this case, the benefit of conservatism costs little in terms of power exactly because true discoveries are rare. Yet when discoveries are more likely, EB behaves more like uncorrected OLS, giving it high power to detect discoveries and suffering little in terms of false discoveries because true positives abound.

The limitations of frequentist MT corrections are well studied in the statistics literature.

Berry and Hochberg (1999) note that “these procedures are very conservative (especially in large families) and have been subjected to criticism for paying too much in terms of power for achieving (conservative) control of selection effects.” The reason is that, while inflating confidence intervals and p -values indeed reduces the discovery of false positives, it also reduces power to detect true positives.

Much of the discussion around MT adjustments in the finance literature fails to consider the loss of power associated with frequentist corrections. But, as Greenland and Hofman (2019) point out, this tradeoff should be a first-order consideration for a researcher navigating multiple tests, and frequentist MT corrections tend to place an implicit cost on false positives that can be unreasonably large. Unlike some medical contexts for example, there is no obvious motivation for asymmetric treatment of false positives and missed positives in factor research. The finance researcher may be willing to accept the risk of a few false discoveries to avoid missing too many true discoveries. In statistics, this is sometimes discussed in terms of an (abstract) cost of Type I versus Type II errors,²⁶ but in finance we can make this cost concrete: We can look at the profit of trading on the discovered factors, where the cost of false discoveries is then the resulting extra risk and money lost (Section 3.3).

2 A New Public Data Set of Global Factors

We study a global dataset with 153 factors in 93 countries. In this section, we provide a brief overview of our data construction. We have posted the code along with extensive documentation detailing every implementation choice that we make for each factor.²⁷

Factors

The set of factors we study is based on the exhaustive list compiled by Hou et al. (2020). They study 202 different characteristic signals from which they build 452 factor portfolios. The

²⁶As Greenland and Robins (1991) point out, “Decision analysis requires, in addition to the likelihood function, a loss function, which indicates the cost of each action under the various possible values for the unknown parameter (benefits would be expressed as negative costs). Construction of a loss function requires one to quantify costs in terms of dollars, lives lost, or some other common scale.”

²⁷It is available at <https://jkpfactors.com/> and at <https://github.com/bkelly-lab/ReplicationCrisis>.

proliferation is due to treating 1, 6, and 12-month holding periods for a given characteristic as different factors, and due to their inclusion of both annual and quarterly updates of some accounting-based factors. In contrast, we focus on a 1-month holding period for all factors, and we only include the version that updates with the most recent accounting data (which could be either annual or quarterly). Lastly, we exclude a small number of factors for which data is not available globally. This gives us a set of 180 feasible global factors. For this set, we exclude factors based on industry or analyst data because they have comparatively short samples.²⁸ This leaves us with 138 factors. Finally, we add 15 factors studied in the literature that were not included in [Hou et al. \(2020\)](#).

For each characteristic, we build the 1-month holding period factor return within each country as follows. First, in each country and month, we sort stocks into characteristic terciles (top/middle/bottom third) with breakpoints based on non-micro stocks in that country.²⁹ For each tercile, we compute its “capped value weight” return, meaning that we weight stocks by their market equity winsorized at the NYSE 80th percentile. This construction ensures that tiny stocks have tiny weights and any one mega stock does not dominate a portfolio, seeking to create tradable, yet balanced, portfolios.³⁰ The factor is then defined as the high-tercile return minus the low-tercile return, corresponding to the excess return of a long-short zero-net-investment strategy. The factor is long (short) the tercile identified by the original paper to have the highest (lowest) expected return.

We scale all factors such that their monthly idiosyncratic volatility is $10\%/\sqrt{12}$ (i.e., 10% annualized), which ensures cross-sectional stationarity and a priori that factors are similar in terms of their information ratio (i.e., appraisal ratio). Finally, we compute each factor’s $\hat{\alpha}^i$ via an OLS regression on a constant and the corresponding region’s market portfolio.

For a factor return to be non-missing, we require that it has at least 5 stocks in each of

²⁸Global industry codes (GICS) are only available from 2000 and I/B/E/S data from the mid-1980’s (but coverage in early years is somewhat sparse).

²⁹Specifically, we start with all non-micro stocks in a country (i.e., larger than NYSE 20th percentile) and sort them into three groups of equal numbers of stocks based on the characteristic, say book-to-market. Then we distribute the micro-cap stocks into the three groups based on the same characteristic breakpoints. This process ensures that the non-micro stocks are distributed equally across portfolios, creating more tradable portfolios.

³⁰For robustness, [Figure C.1](#) of the appendix reports our replication results to using standard, uncapped value weights to construct factors.

the long and short legs. We also require a minimum of 60 non-missing monthly observations for each country-specific factor for inclusion in our sample. When grouping countries into regions (US, developed ex. US, and emerging) we use the MSCI development classification as of January 7th 2021. When aggregating factors across countries, we use capitalization-weighted averages of the country-specific factors. For the developed and emerging market factors, we require that at least three countries have non-missing factor returns.

Clusters

We group factors into clusters using hierarchical agglomerative clustering ([Murtagh and Legendre, 2014](#)). We define the distance between factors as one minus their pairwise correlation and use the linkage criterion of [Ward \(1963\)](#). The correlation is computed based on CAPM-residual returns of US factors signed as in the original paper. Appendix Figure [I.1](#) shows the resulting dendrogram, which illustrates the hierarchical clusters identified by the algorithm. Based on the dendrogram, we choose 13 clusters that demonstrate a high degree of economic and statistical similarity. The cluster names indicate the types of characteristics that dominate each group: Accruals*, Debt Issuance*, Investment*, Leverage*, Low risk, Momentum, Profit Growth, Profitability, Quality, Seasonality, Size*, Skewness*, and Value, where the star (*) indicates that these factors bet against the corresponding characteristic (e.g., accrual factors go long stocks with low accruals while shorting those with high accruals). Appendix Figure [I.2](#) shows that the average within-cluster pairwise correlation is above 0.5 for 9 out of 13 clusters, and Table [J.1](#) provides details on the cluster assignment, sign convention, and original publication source for each factor.

Data and Characteristics

Return data is from CRSP for the US (beginning in 1926) and from Compustat for all other countries (beginning in 1986 for most developed countries).³¹ All accounting data is from Compustat. For international data, all variables are measured in US dollars (based on exchange rates from Compustat) and excess returns are relative to the US treasury bill rate.

³¹Appendix Table [J.3](#) shows start date and other information for all countries included in our dataset.

To alleviate the influence of data errors in the international data, we winsorize returns from Compustat at 0.1% and 99.9% each month.

We restrict our focus to common stocks that are identified by Compustat as the primary security of the underlying firm and assign stocks to countries based on the country of their exchange.³² In the US, we include delisting returns from CRSP. If a delisting return is missing and the delisting is for a performance-based reason, we set the delisting return to -30% following [Shumway \(1997\)](#). In the global data, delisting returns are not available, so all performance-based delistings are assigned a return of -30% .

We build characteristics in a consistent way, that sometimes deviates from the exact implementation used in the original reference. For example, for characteristics that use book equity, we always follow the method in [Fama and French \(1993\)](#). Furthermore, we always use the most recent accounting data, whether annual or quarterly. Quarterly income and cash flow items are aggregated over the previous four quarters to avoid distortions from seasonal effects. We assume that accounting data is available four months after the fiscal period end. When creating valuation ratios, we always use the most recent price data following [Asness and Frazzini \(2013\)](#). Section K in the internet appendix contains a detailed documentation of our data set.

Empirical Bayes Estimation

We estimate the hyperparameters and the posterior alpha distributions of our Bayesian model via EB. Appendix B provides details on the EB methodology and the estimated parameters.

3 Empirical Assessment of Factor Replicability

We now report replication results for our global factor sample. We first present an internal validity analysis by studying US factors over the full sample. Then we analyze external validity in the global cross section and in the time series (post-publication factor returns).

³²Compustat identifies primary securities in the US, Canada and rest of the world. This means that some firms can have up to three securities in our data set. In practice, the vast majority of firms (97%) only have one security in our sample at a given point in time.

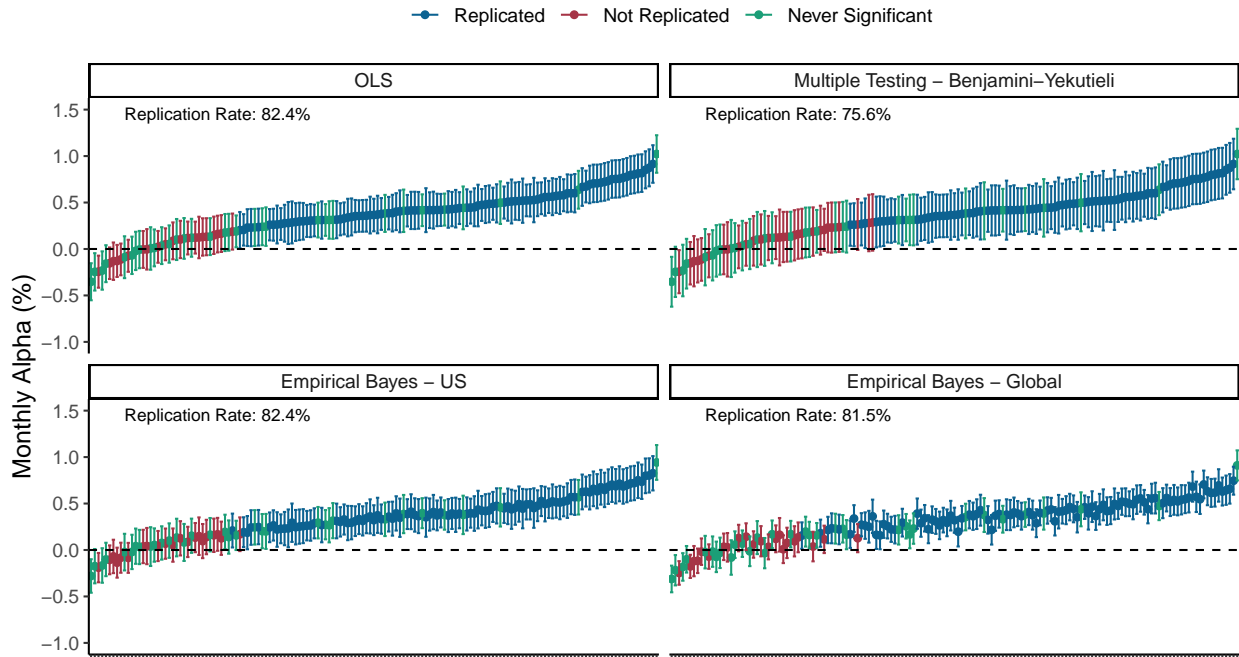


Figure 4: Alpha Distributions for US Factors

Note: The figure reports point estimates and confidence intervals for US factors. The upper left reports OLS estimates. The upper right uses the OLS point estimate but adjusts the confidence interval following the BY procedure. The lower left panel shows our EB posterior confidence intervals using only US data. The lower right continues to show EB results for US factors, but estimates the US factor posterior from global data rather than US-only data. Blue (red) confidence intervals correspond to factors that were significant in the original study in the literature and that we find significant (insignificant) based on the method in each panel. Green intervals correspond to factors that the original study find insignificant or do not evaluate in terms of average return significance. The order of factors is the same in all panels and is arranged from lowest OLS alpha to highest. Table J.2 shows the factor names arranged in the same order.

3.1 Internal Validity

We report full sample performance of US factors in Figure 4. Each panel illustrates the CAPM alpha point estimate of each factor corresponding to the dot at the center of the vertical bars. Vertical bars represent the 95% confidence interval for each estimate. Bar colors differentiate between three types of factors. Blue shows factors that are significant in the original study and remain significant in our full sample. Red shows factors that are significant in the original study but are insignificant in our test. Green shows factors that are not significant in the original study, but are included in the sample of Hou et al. (2020).

The four panels in Figure 4 differ in how the alphas and their confidence intervals are estimated. The upper left panel reports the simple OLS estimate of each alpha, $\hat{\alpha}_{ols}$, and

the 95% confidence intervals based on unadjusted standard errors, $\hat{\alpha}_{\text{ols}} \pm 1.96 \times SE_{\text{ols}}$.³³ The factors are sorted by OLS $\hat{\alpha}$ estimate, and we use this ordering for the other three panels as well. We find that the OLS replication rate is 82.4%, computed as the number of blue factors (98) divided by the sum of red and blue factors (119). Based on OLS tests, factors are highly replicable.

The upper right panel repeats this analysis using the MT adjustment of [Benjamini and Yekutieli \(2001\)](#) (denoted BY), which is advocated by [Harvey et al. \(2016\)](#) and implemented by [Hou et al. \(2020\)](#). This method leaves the OLS point estimate unchanged, but inflates the p -value. We illustrate this visually by widening the alpha confidence interval. Specifically, we find the BY-implied critical value³⁴ in our sample to be a t -statistic of 2.7, and we compute the corresponding confidence interval as $\hat{\alpha}_{\text{ols}} \pm 2.7 \times SE_{\text{ols}}$. We deem a factor as significant according to the BY method if this interval lies entirely above zero. Naturally, this widening of confidence intervals produces a lower replication rate of 75.6%. However, the BY correction does not materially change the OLS-based conclusion that factors appear highly replicable.

The lower left panel is based on our empirical Bayes estimates using the full sample of US factors. For each factor, we use Proposition 4 to compute its posterior mean, $E(\alpha_i | (\hat{\alpha}_j)_{j \text{ any US factor}})$, shown as the dot at the center of the confidence interval. These dots change relative to the OLS estimates, in contrast to BY and other frequentist MT methods that only change the size of the confidence intervals. We also compute the posterior volatility to produce Bayesian confidence intervals, $E(\alpha_i | (\hat{\alpha}_j)_{j \text{ any US factor}}) \pm 1.96 \times \sigma(\alpha_i | (\hat{\alpha}_j)_{j \text{ any US factor}})$. The replication rate based on Bayesian model estimates is 82.4%, larger than BY and, coincidentally, the same as the OLS replication rate. This replication rate has a built-in conservatism from the zero-alpha prior, and it further accounts for the multiplicity of factors because each factor's posterior depends on *all* of the observed evidence in the US (not just own-factor performance).

The lower right panel again reports EB estimates for US factors, but now we allow the

³³We define SE_{ols} as the diagonal of the alpha covariance matrix $\hat{\Sigma}$, which we estimate according to Appendix B.

³⁴We compute the BY-implied critical value as the average of the t -statistic of the factor that is just significant based on BY (the factor with the highest BY-adjusted p -value below 5%) and the t -statistic of the factor that is just insignificant (the factor with the lowest BY-adjusted p -value above 5%).

posterior to depend not just on US data, but on data from all over the world. That is, we compute the posterior mean and variance for each US factor conditional on the alpha estimates for all factors in all regions. The resulting replication rate is 81.5%, which is slightly lower than the EB replication rate using only US data. Some posterior means are reduced due to the fact that some factors have not performed as well outside the US, which affects posterior means for the US through the dependence among global alphas. For example, when the Bayesian model seeks to learn the true alpha of the “US change in book equity” factor, the Bayesian’s conviction regarding positive alpha is reduced by taking into account that the international version of this factor has underperformed the US version.³⁵

To further assess internal validity, we investigate the replication rate for US factors when those factors are constructed from subsamples based on stock size. One of the leading criticisms of factor research replicability is that results are driven by illiquid small stocks whose behavior in large part reflects market frictions and microstructure as opposed to just economic fundamentals or investor preferences. In particular, [Hou et al. \(2020\)](#) argue that they find a low replication rate because they limit the influence of micro-caps. We find that factors demonstrate a high replication rate throughout the size distribution. Panel A of [Figure 5](#) reports replication rates for US size categories shown in the five bars: mega stocks (largest 20% of stocks based on NYSE breakpoints), large stocks (market capitalization between the 80th and 50th percentile of NYSE stocks), small stocks (between the 50th and 20th percentile), micro stocks (between the 20th and 1st percentile), and nano stocks (market capitalization below the 1st percentile).

We see that the EB replication rates in mega and large stock samples are 77.3% and 79.8%, respectively. This is only marginally lower than the overall US sample replication rate of 82.4%, indicating that criticisms of factor replicability based on arguments around stock size or liquidity are largely groundless. For comparison, small, micro, and nano stocks deliver replication rates of 85.7%, 85.7% and 68.1%, respectively.

In [Panel B of Figure 5](#), we report US factor replication rates by theme cluster. 11 out of 13

³⁵To provide a few more details on this example, the US factor based on annual change in book equity (be_gr1a) has a posterior volatility of 0.095% using only US data and 0.077% using global data, leading to a tighter confidence interval with the global data. However, the posterior mean is 0.22% using only US data and 0.13% using global data.




Figure 5: US Replication Rates By Size Group and Theme Cluster

Note: Panel A reports replication rates for US factors formed from subsamples defined by stocks' market capitalization using our EB method. Panel B reports replication rates for US factors in each theme cluster.

themes are replicable with a rate of 50% or better, with the exceptions being the low leverage and size themes. To understand these exceptions, we note that size factors are stronger in emerging markets (bottom panel of Figure G.1) and among micro and nano stocks (bottom panels of Figure G.2). The theoretical foundation of the size effect is a compensation for market illiquidity (Amihud and Mendelson, 1986) and market liquidity risk (Acharya and Pedersen, 2005). Theory predicts that the illiquidity (risk) premium should be the same order of magnitude as the differences in trading costs and these differences are simply much larger in emerging markets and among micro stocks.

Another reason why some factors and themes appear insignificant is that we are not

accounting for other factors. Factors published after 1993 are routinely benchmarked to the Fama-French three-factor model (and, more recently, to the updated five-factor model). Some factors are insignificant in terms of raw return or CAPM alpha, but their alpha becomes significant after controlling for other factors. This indeed explains the lack of replicability for the low leverage theme. While CAPM alphas of low leverage factors are insignificant, we find that it is one of the best performing themes once we account for multiple factors (see Section 3.4 below).

3.2 External Validity

We find a high replication rate in our full-sample analysis, indicating that the large majority of factors are reproducible at least in-sample. We next study the external validity of these results in international data and in post-publication US data.

Global Replication

Figure 6 shows corresponding replication rates around the world. We report replication rates from four testing approaches: (1) OLS with no adjustment; (2) OLS with Benjamini-Yekutieli MT adjustment; (3) the EB posterior conditioning only on factors within a region (“Empirical Bayes – Region”); and (4) EB conditioning on factors in all regions (“Empirical Bayes – All”). Even when using all global data to update the posterior of all factors, the reported Bayesian replication rate applies only to the factors within the stated region.

The first set of bars establishes a baseline by showing replication rates for the US sample, summarizing the results from Figure 4. The next two sets of bars correspond to the developed ex. US sample and the emerging markets sample, respectively.³⁶ Each region factor is a capitalization-weighted average of that factor among countries within a given region, and the replication rate describes the fraction of significant CAPM alphas for these regional factors.

OLS replication rates in developed and emerging markets are generally lower than in

³⁶The developed and emerging samples are defined by the MSCI development classification and include 23 and 27 countries, respectively. The remaining 43 countries in our sample that are classified as neither developed nor emerging by MSCI do not appear in our developed and emerging region portfolios, but they are included in the “world” versions of our factor portfolios.




Figure 6: Replication Rates in Global Data

Note: We report replication rates for factors in three global regions (US, developed ex. US, and emerging) and for the world as a whole. A factor in a given region is the capitalization-weighted average factor for countries in that region. We report OLS replication rates with no adjustment and with Benjamini-Yekutieli multiple testing adjustment. We also report replication rates based on the empirical Bayes posterior. We consider two EB methods. In both methods, the replication rate refers only to factors within the region of interest, but the posterior is computed by conditioning either on data from that region alone (“Empirical Bayes – Region”) or on the full global sample (“Empirical Bayes – All”). We deem a factor successfully replicated if its 95% confidence interval excludes zero for a given method.

the US, and the frequentist Benjamini-Yekutieli correction has an especially large negative impact on replication rate. This is a case in which the Bayesian approach to MT is especially powerful. Even though the alphas of all regions are shrunk toward zero, the global information set helps EB achieve a high degree of precision, narrowing the posterior distribution around the shrunk point estimate. We can see this in increments. First, the EB replication rate using region-specific data (“Empirical Bayes – Region” in the figure) is just below the OLS replication rate but much higher than the Benjamini-Yekutieli rate. When the posterior leverages global data (“Empirical Bayes – All” in the figure), the replication rate is higher still, reflecting the benefits of sharing information across regions, as recommended by the dependence among alphas in the hierarchical model.

Finally, we use the global model to compute, for each factor, the capitalization-weighted




Figure 7: US Factor Alphas Versus World Ex. US

Note: The figure compares OLS alphas for US factors versus their international counterpart. Each world ex. US factor is a capitalization-weighted average of the factor in all other countries of our sample. Blue points correspond to factors that were significant in the original study in the literature, while red points are those for which the original paper did not find a significant effect (or did not study the factor in terms of average return significance). The dotted line is the 45° line. The figure also reports a regression of world ex. US alpha on US alpha.

average alpha across all countries in our sample (“World” in the figure). Using data from around the world, we find a Bayesian replication rate of 82.4%.

Why do international OLS replication rates differ from the US? This is due primarily to the the fact that foreign markets have shorter time samples. Point estimates are similar in magnitude for the US and international data. Figure 7 shows the alpha of each US factor against the alpha of the corresponding factor for the world ex. US universe. The data cloud aligns closely with the 45° line, demonstrating the close similarity of alpha magnitudes in the two samples. But shorter international samples widen confidence intervals, and this is the primary driver of the drop in OLS replication rates outside the US.




Figure 8: In-Sample versus Out-of-Sample Alphas for US Factors

Note: The figure plots OLS alphas for US factors during the in-sample period (i.e., the period studied in the original publication) versus out-of-sample alphas. In Panel A, out-of-sample is the time period before the in-sample period. In Panel B, out-of-sample is the time period before the in-sample period. In Panel C, out-of-sample includes both the time period before and after the in-sample period. We require at least five years of out-of-sample data for a factor to be included, amounting to 102, 115 and 119 factors in panel A, B and C. The figure also reports feasible GLS estimates of out-of-sample alphas on in-sample alphas. To implement feasible GLS, we assume that the error variance-covariance matrix is proportional to the full-sample CAPM residual variance-covariance matrix, $\hat{\Sigma}/T$, described in section B. The dotted line is the 45° line.

Time Series Out-of-Sample Evidence

McLean and Pontiff (2016) document the intriguing fact that, following publication, factor performance tends to decay. They estimate an average post-publication decline of 58% in factor returns. In our data, the average in-sample alpha is 0.49% per month and the average out-of-sample alpha is 0.26% when looking post-original sample, implying a decline of 47%.

We gain further economic insight by looking at these findings cross-sectionally. Figure 8 makes a cross-sectional comparison of the in-sample and out-of-sample alphas of our US factors. The in-sample period is the sample studied in the original reference. The out-of-sample period in Panel A is the time period before the start of in-sample period, while in Panel B it is the period following the in-sample period. Panel C defines out-of-sample as the combined data from the periods before and after the originally studied sample. We find that 82.6% of the US factors that were significant in the original publication also have positive returns in the pre-original sample, 83.3% are positive in the post-original sample, and 87.4% are positive in the combined out-of-sample period. When we regress out-of-sample alphas on in-sample alphas using GLS, we find a slope coefficient of 0.57, 0.26, and 0.35 in Panels

A, B, and C, respectively. The slopes are highly significant (ranging from $t = 3.5$ to $t = 5.3$) indicating that in-sample alphas contain something “real” rather than being the outcome of pure data mining, as factors that performed better in-sample also tend to perform better out-of-sample.

The significantly positive slope allows us to reject the hypothesis of “pure alpha-hacking,” which would imply a slope of zero, as seen in Proposition 1. Further, the regression intercept is positive, while alpha-hacking of the form studied in Proposition 1 would imply a negative intercept.

That the slope coefficient is positive and less than one is consistent with basic Bayesian logic of equation (4). As we emphasize in Section 1, a Bayesian would expect at least some attenuation in out-of-sample performance. This is because the published studies report the OLS estimate, while Bayesian beliefs shrink the OLS estimate toward the zero-alpha prior. More specifically, with no alpha hacking or arbitrage, the Bayesian expects a slope of approximately 0.9 using equation (5) and our EB hyperparameters (see appendix Table B.1).³⁷ Hence, the slope coefficients in Figure 8 are too low relative to this Bayesian benchmark. In addition to the moderate slope, there is evidence that the dots in Figure 8 have a concave shape (as seen more clearly in appendix Figure D.2). These results indicate that, while we can rule out pure alpha-hacking (or p -hacking), there is some evidence that the highest in-sample alphas may either be data-mined or arbitrated down.

From the Bayesian perspective, another interesting evaluation of time series external validity is to ask whether the new information contained in out-of-sample data moves the posterior alpha toward zero or not. Imagine a Bayesian observing the arrival of factor data in real time. As new data arrives, she updates her beliefs for all factors based on the information in the full cross section of factor data. In the top panel of Figure 9, we show how the Bayesian’s posterior of the average alpha would have evolved in real time.³⁸ We focus on all the World factors that are available since at least 1955 and significant in the original

³⁷The slope is $\kappa = 1/(1 + \sigma^2/(T\tau^2)) = 0.9$, where $\sigma^2 = 10\%^2/12$, the average in-sample period length is $T = 420$ months, and $\tau^2 = \tau_c^2 + \tau_w^2 = (0.35\%)^2 + (0.21\%)^2 = (0.41\%)^2$.

³⁸Here we keep τ_c and τ_w fixed at their full-sample values of 0.37% and 0.23% to mimic the idea of given decision maker who starts with a given prior and updates this view based on new data, while keeping the prior fixed. Figure D.3 shows that the figure is almost the same with rolling estimates of τ_c and τ_w , and Figure D.4 shows that this consistency arises because the rolling estimates are relatively stable.




Figure 9: World Factor Alpha Posterior Distribution Over Time

Note: The top panel reports the CAPM alpha and 95% posterior confidence interval for an equal-weighted portfolio of World factors based on EB posteriors re-estimated in December each year. That is, each blue dot is $E(\frac{1}{N} \sum_i \alpha^i | \text{data until time } t)$ and the vertical lines are ± 2 times the posterior volatility. Triangles show average OLS alpha at each point in time, $\frac{1}{N} \sum_i \hat{\alpha}_{t_0, t}^i$, estimated using data through date t . The bottom panel reports the average monthly alpha for all factors in a rolling 5-year window. The results are based on factors found to be significant in the original paper with data available since 1955.

paper. Starting in 1960, we re-estimate the hierarchical model using the empirical Bayes estimator in December of each year. The plot shows the CAPM alpha and corresponding 95% confidence interval of an equal-weighted portfolio of the available factors. The posterior mean alpha becomes relatively stable from the mid 1980s, around 0.4% per month. And, as data evidence has accumulated over time, the confidence interval narrows by a third, from about 0.16% wide in 1960 to 0.10% in 2020.

To understand the posterior alpha, Figure 9 also shows the average OLS alpha as triangles and the bottom panel in Figure 9 reports the rolling 5-year average monthly alpha among all these factors. We see that the EB posterior is below the OLS estimate, which occurs because the Bayesian posterior is shrunk toward the zero prior. Naturally, periods of good performance increase the posterior mean as well as the OLS estimate, and vice versa for poor performance. Over time, the OLS estimate moves nearer to the Bayesian posterior mean.

To further understand why the posterior alpha is relatively stable with a tightening confidence interval, consider the following simple example. Suppose a researcher has $T = 10$ years of data for factors with an OLS alpha estimate of $\hat{\alpha} = 10\%$ with standard error σ/\sqrt{T} . Further, assume their zero-alpha prior is equally as informative as their 10-year sample (i.e., $\tau = \sigma/\sqrt{T}$). Then the shrinkage factor is $\kappa = 1/2$ using equation (5). So, after observing the first ten years with $\hat{\alpha} = 10\%$, the Bayesian expects a future alpha of $E(\alpha|\hat{\alpha}) = 5\%$ (equation (4)). What happens if this Bayesian belief is confirmed by additional data, namely that the factor realizes an alpha of 5% over the next 10 years? In this case, the full-sample OLS of alpha is $\hat{\alpha} = 7.5\%$, but now the shrinkage factor becomes $\kappa = 2/3$ because the sample length doubles, $T = 20$. This results in a posterior alpha of $E(\alpha|\hat{\alpha}) = 7.5\% \cdot 2/3 = 5\%$. Naturally, when beliefs are confirmed by additional data, the posterior mean does not change. Nevertheless, we learn something from the additional data, because our conviction increases as the posterior variance is reduced. If $\sigma = 0.1$, the posterior volatility $\sqrt{\text{Var}(\alpha|\hat{\alpha})} = \sigma\sqrt{\frac{\kappa}{T}}$ goes from 2.2% with 10 years of data to 1.8% with 20 years of data, and the confidence interval, $[E(\alpha|\hat{\alpha}) \pm 2\sqrt{\text{Var}(\alpha|\hat{\alpha})}]$, is reduced from [0.5%, 9.5%] to [1.3%, 8.7%].

3.3 Bayesian Multiple Testing

A great advantage of Bayesian methods for tackling challenges in multiple testing is that, from the posterior distribution, we can make explicit probability calculations for essentially any inferential question. We simulate from our EB posterior to investigate the false discovery and family-wise error rates among the set of global factors that were significant in the original study. We define a false discovery as a factor where we claim that the alpha is positive, but where the true alpha is negative.³⁹

First, based on Proposition 5, we calculate the Bayesian FDR in our sample as the average posterior probability of a false discovery, p -null, among all discoveries. We find that $\text{FDR}^{\text{Bayes}} = 0.1\%$, meaning that we expect roughly one discovery in 1000 to be a false positive given our Bayesian hierarchical model estimates. The posterior standard error for

³⁹In particular, we define a discovery as a factor for which the posterior probability of the true alpha being negative is less than 2.5%. With this definition, we start with 153 world factors, then focus on the 119 factors that were significant in the original studies, and, out of these, 98 are considered discoveries.

$\text{FDR}^{\text{Bayes}}$ is 0.3% with a confidence interval of [0,1%]. In other words, the model generates a highly conservative MT adjustment in the sense that once a factor is found to be significant, we can be relatively confident that the effect is genuine.

We can also use the posterior to make other inference calculations. We compute the FWER, which we define as the probability of at least one false discovery. We simulate 1,000,000 draws of the vector of alphas that were deemed to be discoveries from the EB posterior and compute

$$\text{FWER}^{\text{Bayes}} = \frac{1}{1,000,000} \sum_{s=1}^{1,000,000} 1_{\{n_s \geq 1\}} = 5.5\%$$

where n_s is the number of false discoveries in simulation s . In other words, the probability of at least one alpha having the wrong sign is 5.5%. The $\text{FWER}^{\text{Bayes}}$ is naturally much higher than the $\text{FDR}^{\text{Bayes}}$ given the extreme conservatism built into the FWER's definition of false discovery. Whether it is too high is subjective. A nice aspect of our approach is that a researcher can control the $\text{FWER}^{\text{Bayes}}$ as desired. For example, using a t -statistic threshold of 2.78 rather than 1.96 leads to $\text{FWER}^{\text{Bayes}} = 0.8\%$.

From the posterior, we can also compute the expected fraction of discovered factors that are “true,” which is in general different than the replication rate. The replication rate is the fraction of factors having $E(\alpha_i|\text{data})/\sigma(\alpha_i|\text{data}) > 1.96$, while the expected fraction of true factors is $\frac{1}{n} \sum_i E(1_{\alpha_i > 0}|\text{data}) = \frac{1}{n} \sum_i Pr(\alpha_i > 0|\text{data})$. The replication rate gives a conservative take on the number of true factors—the expected fraction of true factors is typically higher than the replication rate. To understand this conservatism, consider an example in which all factors have a 90% posterior probability of being true. These would all individually be counted as “not replicated,” but they would contribute to a high expected fraction of true factors. Indeed, we estimate that the expected fraction of factors with truly positive alphas is 94% (with a posterior standard error of 1.3%), notably higher than our estimated replication rate.

Economic Benefits of More Powerful Tests

MT adjustments should ultimately be evaluated by whether they lead to better decisions. It is important to balance the relative costs of false positives versus false negatives, and the appropriate tradeoff depends on the context of the problem (Greenland and Hofman, 2019). We apply this general principle in our context by directly measuring costs in terms of investment performance.

Specifically, we can compute the difference in out-of-sample investment performance from investing using factors chosen with different methods. We compare two alternatives. One is the BY decision rule advocated by Harvey et al. (2016), which is a frequentist MT method that successfully controls false discoveries relative to OLS, but in doing so sacrifices power (the ability to detect true positives). The second alternative is our EB method, whose false discovery control typically lies somewhere between BY and unadjusted OLS. EB uses the data sample itself to decide whether its discoveries should behave more similarly to BY or to unadjusted OLS.

For investors, the optimal decision rule is the one that leads to the best performance out-of-sample. For the most part, the set of discovered factors for BY and EB coincide. It is only in marginal cases where they disagree which, in our sample, occurs when EB makes a discovery that BY deems insignificant. Therefore, to evaluate MT approaches in economic terms, we track the out-of-sample performance of factors included by EB but excluded by BY. If the performance of these is negative on average, then the BY correction is warranted and preferred by the investor.

We find that the out-of-sample performance of factors discovered by EB but not BY is positive on average and highly significant. The alpha for these marginal cases is 0.35% per month among US factors ($t = 5.1$).⁴⁰ This estimate suggests that the BY decision rule is too conservative. An investor using the rule would fail to invest in factors that subsequently have a high out-of-sample return.

Another way to see that the BY decision rule is too conservative comes from the connec-

⁴⁰For the developed ex. US sample, the monthly alpha for marginal cases is 0.24% per month ($t = 5.3$), and for the emerging sample it is 0.27% ($t = 3.7$), in favor of the EB decision rule. Appendix Table E.1 reports additional details for this analysis.

tion between Sharpe ratio and t -statistics: $t = \text{SR}\sqrt{T}$. If we have a factor with an annual Sharpe ratio of 0.5, an investor using the 1.96 cutoff would in expectation invest in the factor after 15 years. An investor using the 2.78 cutoff, would not start investing until observing the factor for 31 years.

Addressing Unobserved Factors, Publication Bias, and other Biases

A potential concern with our replication rate is that the set of factors that make it into the literature is a selected sample. In particular, researchers may have tried many different factors, some of which are observed in the literature, while others are unobserved because they never got published. Unobserved factors may have worse average performance if poor performance makes publication more difficult or less desirable. Alternatively, unobserved factors could have strong performance if people chose to trade on them in secret rather than publishing them. Either way, we next show how unobserved factors can be addressed in our framework.

The key insight is that the performance of factors across the universe of observed and unobserved factors is captured in our prior parameters τ_c, τ_w . Indeed, large values of these priors correspond to a large dispersion of alphas (that is, a lot of large alphas “out there”) while small values means that most true alphas are close to zero. Therefore, smaller τ 's lead to a stronger shrinkage toward zero for our posterior alphas, leading to fewer factor “discoveries” and a lower replication rate. Figure 10 shows how our estimated replication rate depends on the most important prior parameter, τ_c , based on the τ_w that we estimated from the data.⁴¹

In Figure 10, we show how the replication rate varies with τ_c in precise quantitative terms. Note that while the replication rate indeed rises with τ_c , the differences are small in magnitude across a large range of τ_c values, demonstrating robustness of our conclusions about replicability.

This stable replication rate in Figure 10 also suggests that the replication rate among the observed factors would be similar even if we had observed the unobserved factors. The figure highlights several key values of τ_c : Both the value of τ_c that we estimated from the

⁴¹Figure F.1 in the appendix shows that the results are robust to alternative values of τ_w .




Figure 10: Replication Rate with Prior Estimated in Light of Unobserved Factors

Note: The figure shows how the replication rate in the US varies when changing the τ_c parameter. The τ_w parameter is fixed at the estimate value of 0.21%. The dotted line shows our replication rate of 82.4%. The green square, highlights the value estimated in the data $\tau_c = 0.35\%$. The red triangle and the blue circle highlights values that are found by estimating the empirical Bayes model according to assumptions about unobserved factors from Harvey et al. (2016). The values are $\tau_c = 0.28\%$ in the baseline scenario and $\tau_c = 0.20\%$ in the conservative scenario. A description of this approach can be found in Appendix F.

observed data (as explained in Appendix B) and values that adjust for unobserved data in different ways.

We adjust τ_c for unobserved factors as follows. We simulate a data set that proxies for the full set of factors in the population (including those unobserved), and then estimate the τ 's that match this sample. One set of simulations is constructed to match the baseline scenario of Harvey et al. (2016) (Table 5.A, row 1), which estimates that researchers have tried $M = 1,297$ factors, of which 39.6% of have zero alpha and the rest have a Sharpe ratio of 0.44. We also consider the more conservative scenario of Harvey et al. (2016) (Table 5.B, row 1), which implies that researchers have tried $M = 2,458$ factors, of which 68.3% have zero alpha. Appendix F has more details on these simulations. The result, as seen in Figure 10, is that values of τ_c that correspond to these scenarios from Harvey et al. (2016) still lead to a conclusion of a high replication rate in our factor universe. The replication rate




Figure 11: World Alpha Posterior By Factor and Cluster

Note: The figure reports the EB posterior 95% confidence interval for the true alpha of a world factor create as a capitalization weighted average of all country specific factors in our dataset. We only include factors that the original paper finds significant.

is 81.5%, and 79.8% for the prior hyperparameters implied by the baseline and conservative scenario respectively.

A closely related bias is that factors may suffer from alpha-hacking as discussed in Section 1.1 (Proposition 1), which makes realized in-sample factor returns too high. To account for this bias, we estimate the prior hyper-parameters using only out-of-sample data. These estimated values are $\tau_c = 0.27\%$ and $\tau_w = 0.22\%$. These hyper-parameters are similar to those implied by the baseline scenario of Harvey et al. (2016) as seen in Figure 10. With these hyper-parameters, the replication rate is 81.5%.

3.4 Economic Significance of Factors

Which factors (and which themes) are the most impactful anomalies in economic terms? We investigate this question by identifying which factors matter most from an investment performance standpoint.

Figure 11 shows the alpha confidence intervals for all world factors, sorted by the median

posterior alpha within clusters. This illustration is similar to Figure 4, but now we focus on the world instead of the US factors, and here we sort factors into clusters. We also focus on factors that the original studies conclude are significant. We see that world factor alphas tend to be economically large, often above 0.3% per month, and tend to be highly significant, in most clusters. The exception is the low leverage cluster, where we also saw a low replication rate in preceding analyses.

By Region and By Size

We next consider which factors are most economically important across global regions and across stock size groups. In Panel A of Figure 12, we construct factors using only stocks in the five size subsamples presented earlier in Figure 5; namely mega, large, small, micro, and nano stock samples. For each sample, we calculate cluster-level alphas as the equal-weighted average alpha of rank-weighted factors within the cluster.⁴² We see, perhaps surprisingly, that the ordering and magnitude of alphas is broadly similar across size groups. The Spearman rank correlation of alphas for mega caps versus micro caps is 73%. Only the nano stock sample, defined as stocks below the 1st percentile of the NYSE size distribution (which amounted to 458 out of 4356 stocks in the US at the end of 2020), exhibits notable deviation from the other groups. The Spearman rank correlation between alphas of mega caps and nano caps is 36%.

Panel B of Figure 12 shows cluster-level alphas across regions. Again, we find consistency in alphas across the globe, with the obvious standout being the size theme, which is much more important in emerging markets than in developed markets. US factor alphas share a 62% Spearman correlation with the developed ex. US sample, and a 43% correlation with the emerging markets sample.

Controlling for Other Themes

We have focused so far on whether factors (or clusters) possess significant positive alpha relative to the market. The limitation of studying factors in terms of CAPM alpha is that it

⁴²Rank-weighting is similar to equal-weighting and used here to illustrate the performance of typical stocks in each size group. See equation (1) in [Asness et al. \(2013\)](#).

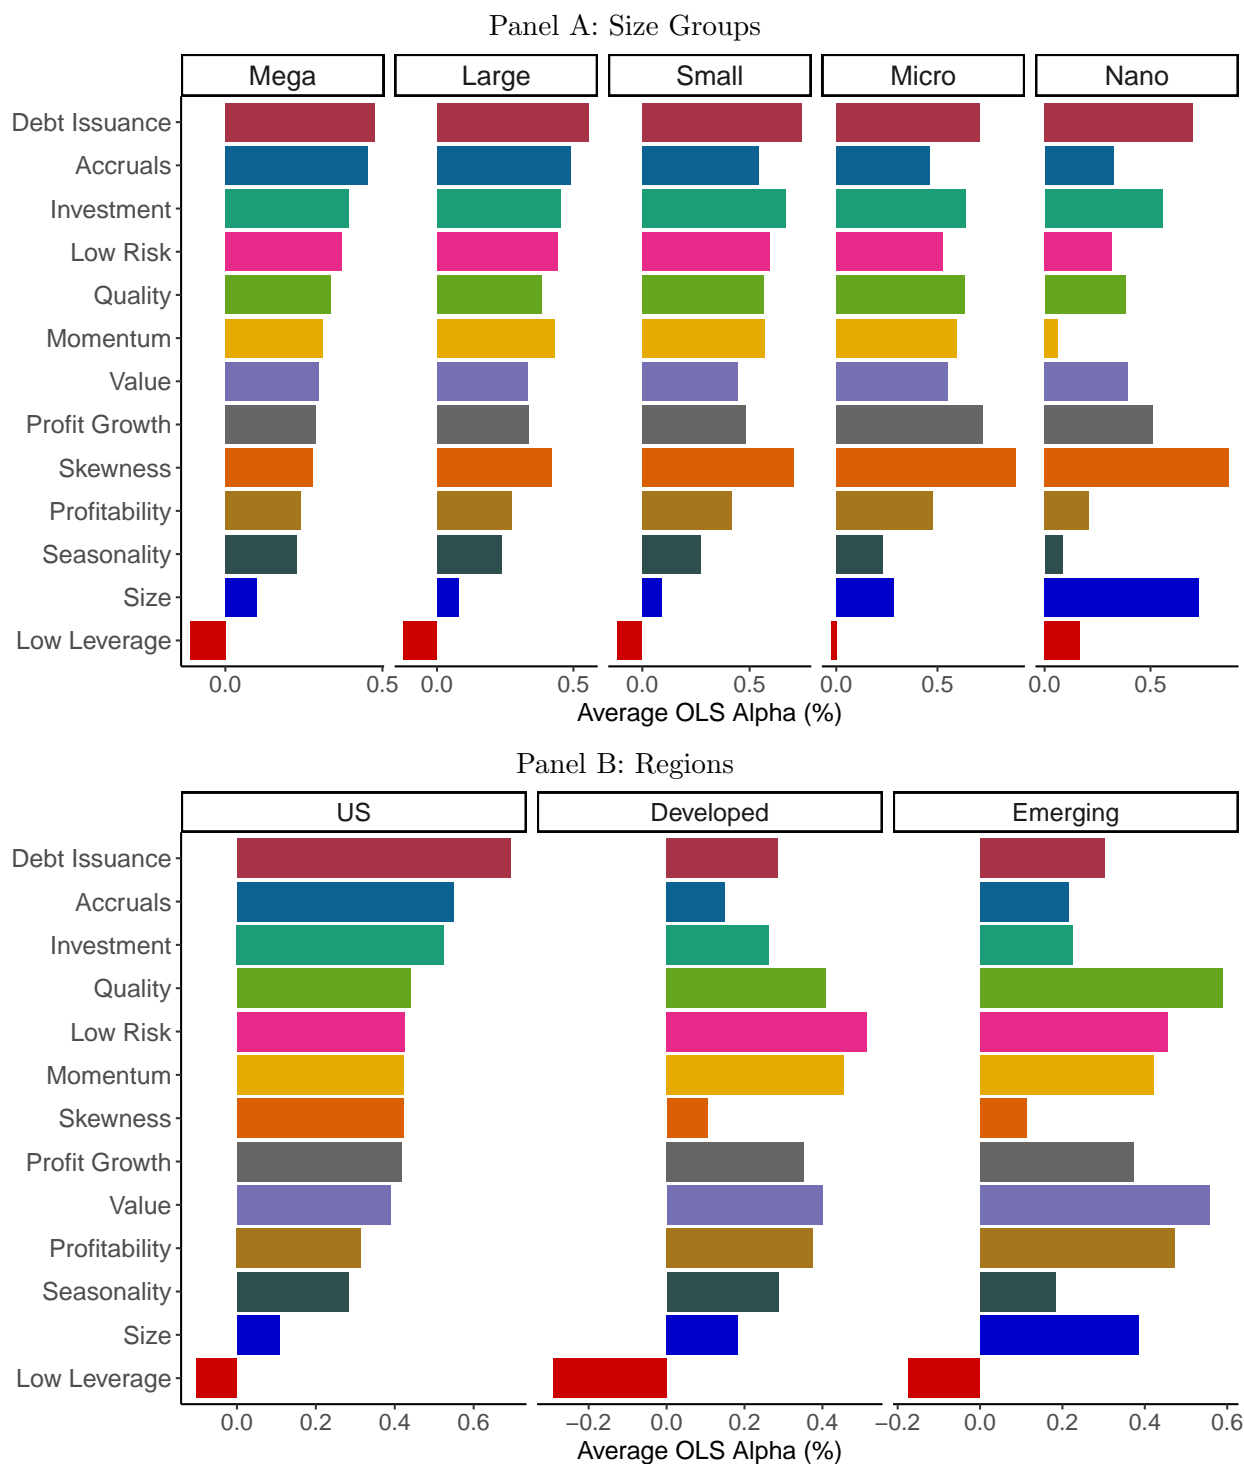


Figure 12: Alphas By Geographic Region and Stock Size Group

Note: The figure reports average cluster-level alphas for factors formed from subsamples defined by different stock market capitalization groups (Panel A) and regions (Panel B).




Figure 13: Tangency Portfolio Weights

Note: The return are from the US portfolios. We compute the cluster return as the equal weighted return of all factors with data available at a given point in time. We further add the US market return. We estimate the tangency weights following the method of Britten-Jones (1999) with a non-negativity constraint. The error bars are the 90% confidence intervals based on 10,000 bootstrap samples and the percentile method. The data starts in 1952 to ensure that all cluster have non-missing observations.

does not control for duplicate behavior other than through the market factor. Economically important factors are those that have large impact on an investor’s overall portfolio, and this requires understanding which clusters contribute alpha while controlling for all others.

To this end, we estimate cluster weights in a tangency portfolio that invests jointly in all cluster-level portfolios. We test the significance of the estimated weights using the method of Britten-Jones (1999). In addition to our 13 cluster-level factors, we also include the market portfolio as a way of benchmarking factors to the CAPM null. Lastly, we constrain all weights to be non-negative (because we have signed the factors to have positive expected returns according to the findings of the original studies).

Figure 13 reports the estimated tangency portfolio weights and their 90% bootstrap confidence intervals. When a factor has a significant weight in the tangency portfolio, it means that it matters for an investor, even controlling for all the other factors. We see that all but three clusters are significant in this sense. We also see that conclusions about




Figure 14: The Evolution of the Tangency Sharpe ratio

Note: The top panel shows the Sharpe ratio on the ex-post tangency portfolio. A factor is included in the tangency portfolio only after the end of the sample in which the factor was studied in the original publication (and we only include factors that were found to be significant in the original paper). We highlight selected factors that significantly improve the optimal portfolio, starting with the market portfolio. We use the longest available balanced US sample, 1972–2020 (that is, when all factors are available).

cluster importance change when clusters are studied jointly. For example, value factors become stronger when controlling for other effects because of their hedging benefits relative to momentum, quality, and low leverage. More surprisingly, the low leverage cluster becomes one of the most heavily weighted clusters, in large part due to its ability to hedge value and low risk factors. The hedging performance of value and low leverage clusters is clearly discernible in Appendix table I.2, which shows the average pairwise correlations among factors within and across clusters.⁴³ Appendix H provides further performance attribution of the tangency portfolio at the factor level.⁴⁴

Evolution of Finance Factor Research

The number of published factors has increased over time as seen in the bottom panel of Figure 14. But, to what extent have these new factors continued to add new insights versus repackaging existing information?

⁴³Appendix Tables H.1 and H.2 show how tangency portfolio weights vary by region and by size group.

⁴⁴Figure H.3 shows the performance of each cluster in combination with the market portfolio, figure H.4 shows how the optimal portfolio changes when one cluster is excluded and figure H.6 shows the importance of each factor for the optimal portfolio.

To address this question, we consider how the optimal risk-return tradeoff has evolved over time as factors have been discovered. Specifically, Figure 14 computes the monthly Sharpe ratio of the ex-post tangency portfolio that only invests in factors discovered by a certain point in time.⁴⁵ The starting point (on the left) of the analysis is the 0.13 Sharpe ratio of the market portfolio in the US sample 1972-2020 when all factors are available. The ending point (on the right) is the 0.80 Sharpe ratio of the tangency portfolio that invests the optimal weights across all factors over the same US sample period.⁴⁶ In between, we see how the Sharpe ratio of the tangency portfolio has evolved as factors have been discovered. The improvement is gradual over time, but we also see occasional large increases when researchers have discovered especially impactful factors (usually corresponding to new themes in our classification scheme). An example is the operating accruals factor proposed by Sloan (1996), which increased the tangency Sharpe ratio from 0.43 to 0.56. More recently, the seasonality factors of Heston and Sadka (2008) increase the Sharpe ratio from 0.65 to 0.74.

4 Conclusion: Finance Research Posterior

We introduce a hierarchical Bayesian model of alphas that emphasizes the joint behavior of factors and provides a more powerful multiple testing adjustment than common frequentist methods. Based on this framework, we re-visit the evidence on replicability in factor research and come to substantially different conclusions versus the prior literature. We find that US equity factors have a high degree of internal validity in the sense that over 80% of factors remain significant after modifications in factor construction that make all factors consistent, more implementable, while still capturing the original signal (Hamermesh, 2007) and after accounting for multiple testing concerns (Harvey et al., 2016; Harvey, 2017).

We also provide new evidence demonstrating a high degree of external validity in factor research. In particular, we find highly similar qualitative and quantitative behavior in a

⁴⁵We estimate tangency portfolio weights following the method of Pedersen et al. (2021), which offers a sensible approach to mean-variance optimization for high dimensional data. Estimation details are provided in Appendix H.

⁴⁶The high Sharpe ratio partly reflects the fact that we are doing an in-sample optimization. If we instead do a pseudo out-of-sample analysis via cross-validation, we find a monthly Sharpe ratio of 0.56.

large sample of 153 factors across 93 countries as we find in the US. We also show that, within the US, factors exhibit a high degree of consistency between their published in-sample results and out-of-sample data not considered in the original studies. We show that some out-of-sample factor decay is to be expected in light of Bayesian posteriors based on publication evidence. Therefore, the new evidence from post-publication data largely confirms the Bayesian's beliefs, which has led to relatively stable Bayesian alpha estimates over time.

In addition to providing a powerful tool for replication, our Bayesian framework has several additional applications. For example, the model can be used to correctly interpret out-of-sample evidence, look for evidence of alpha-hacking, compute the expected number of false discoveries and other relevant statistics based on the posterior, analyze portfolio choice taking into account both estimation uncertainty and return volatility, and evaluate asset pricing models.

Finally, the code, data, and meticulous documentation for our analysis are available online. Our large global factor data set and the underlying stock-level characteristics are easily accessible to researchers by using our publicly available code and its direct link to WRDS. Our database will be updated regularly with new data releases and code improvements. We hope that our methodology and data will help promote credible finance research.

References

- Acharya, V. and L. H. Pedersen (2005). Asset pricing with liquidity risk. *Journal of Financial Economics* 77, 375–410.
- Amihud, Y. and H. Mendelson (1986). Asset pricing and the bid-ask spread. *Journal of Financial Economics* 17, 223–249.
- Asness, C. and A. Frazzini (2013). The devil in HML’s details. *The Journal of Portfolio Management* 39(4), 49–68.
- Asness, C., T. Moskowitz, and L. H. Pedersen (2013). Value and momentum everywhere. *The Journal of Finance* 68(3), 929–985.
- Benjamini, Y. and Y. Hochberg (1995). Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal statistical society: series B (Methodological)* 57(1), 289–300.
- Benjamini, Y. and D. Yekutieli (2001). The control of the false discovery rate in multiple testing under dependency. *The Annals of Statistics* 29(4), 1165–1188.
- Berry, D. A. and Y. Hochberg (1999). Bayesian perspectives on multiple comparisons. *Journal of Statistical Planning and Inference* 82(1-2), 215–227.
- Bettis, R. A. (2012). The search for asterisks: Compromised statistical tests and flawed theories. *Strategic Management Journal* 33(1), 108–113.
- Britten-Jones, M. (1999). The sampling error in estimates of mean-variance efficient portfolio weights. *The Journal of Finance* 54(2), 655–671.
- Bryzgalova, S., J. Huang, and C. Julliard (2019). Bayesian solutions for the factor zoo: We just ran two quadrillion models. *Available at SSRN*.
- Chen, A. Y. (2020). The limits of p-hacking: Some thought experiments. *The Journal of Finance, forthcoming*.
- Chen, A. Y. and T. Zimmermann (2020a). Open source cross-sectional asset pricing. *Working paper, Board of Governors of the Federal Reserve System*.
- Chen, A. Y. and T. Zimmermann (2020b). Publication bias and the cross-section of stock returns. *The Review of Asset Pricing Studies* 10(2), 249–289.
- Chinco, A., A. Neuhierl, and M. Weber (2020). Estimating the anomaly base rate. *Journal of Financial Economics*.
- Chordia, T., A. Goyal, and A. Saretto (2020). Anomalies and false rejections. *The Review of Financial Studies* 33(5), 2134–2179.
- Cochrane, J. H. (2011). Presidential address: Discount rates. *Journal of Finance* 66, 1047–1108.
- Efron, B. (2007). Size, power and false discovery rates. *The Annals of Statistics* 35(4), 1351–1377.
- Efron, B. (2012). *Large-scale inference: empirical Bayes methods for estimation, testing, and prediction*, Volume 1. Cambridge University Press.
- Efron, B. and R. Tibshirani (2002). Empirical bayes methods and false discovery rates for microarrays. *Genetic epidemiology* 23(1), 70–86.

- Elton, E. J., M. J. Gruber, and J. Spitzer (2006). Improved estimates of correlation coefficients and their impact on optimum portfolios. *European Financial Management* 12(3), 303–318.
- Engle, R. and B. Kelly (2012). Dynamic equicorrelation. *Journal of Business & Economic Statistics* 30(2), 212–228.
- Fama, E. F. and K. R. French (1993). Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics* 33(1), 3–56.
- Feng, G., S. Giglio, and D. Xiu (2020). Taming the factor zoo: A test of new factors. *The Journal of Finance* 75(3), 1327–1370.
- Frazzini, A. and L. H. Pedersen (2014). Betting against beta. *Journal of Financial Economics* 111(1), 1 – 25.
- Gelman, A. (2016, Aug). Bayesian inference completely solves the multiple comparisons problem.
- Gelman, A., J. B. Carlin, H. S. Stern, D. B. Dunson, A. Vehtari, and D. B. Rubin (2013). *Bayesian data analysis*. CRC press.
- Gelman, A., J. Hill, and M. Yajima (2012). Why we (usually) don't have to worry about multiple comparisons. *Journal of Research on Educational Effectiveness* 5(2), 189–211.
- Green, J., J. R. Hand, and X. F. Zhang (2017). The characteristics that provide independent information about average us monthly stock returns. *The Review of Financial Studies* 30(12), 4389–4436.
- Greenland, S. and A. Hofman (2019). Multiple comparisons controversies are about context and costs, not frequentism versus bayesianism. *European journal of epidemiology* 34(9), 801–808.
- Greenland, S. and J. M. Robins (1991). Empirical-bayes adjustments for multiple comparisons are sometimes useful. *Epidemiology*, 244–251.
- Hamermesh, D. S. (2007). Replication in economics. *Canadian Journal of Economics/Revue canadienne d'économique* 40(3), 715–733.
- Harvey, C. R. (2017). Presidential address: The scientific outlook in financial economics. *The Journal of Finance* 72(4), 1399–1440.
- Harvey, C. R., Y. Liu, and H. Zhu (2016). ... and the cross-section of expected returns. *The Review of Financial Studies* 29(1), 5–68.
- Heston, S. L. and R. Sadka (2008). Seasonality in the cross-section of stock returns. *Journal of Financial Economics* 87(2), 418–445.
- Hou, K., C. Xue, and L. Zhang (2020). Replicating anomalies. *The Review of Financial Studies* 33(5), 2019–2133.
- Ilmanen, A., R. Israel, T. J. Moskowitz, R. Lee, and A. K. Thapar (2021). How do factor premia vary over time? a century of evidence. *Journal of Investment Management, Forthcoming*.
- Ioannidis, J. P. (2005). Why most published research findings are false. *PLoS medicine* 2(8), e124.
- Jacobs, H. and S. Müller (2020). Anomalies across the globe: Once public, no longer existent? *Journal of Financial Economics* 135(1), 213–230.

- Kelly, B. T., S. Pruitt, and Y. Su (2019). Characteristics are covariances: A unified model of risk and return. *Journal of Financial Economics* 134(3), 501–524.
- Koijen, R. S., T. J. Moskowitz, L. H. Pedersen, and E. B. Vrugt (2018). Carry. *Journal of Financial Economics* 127(2), 197–225.
- Kozak, S., S. Nagel, and S. Santosh (2020). Shrinking the cross-section. *Journal of Financial Economics* 135(2), 271–292.
- Linnainmaa, J. T. and M. R. Roberts (2018). The history of the cross-section of stock returns. *The Review of Financial Studies* 31(7), 2606–2649.
- Maniadis, Z., F. Tufano, and J. A. List (2017). To replicate or not to replicate? exploring reproducibility in economics through the lens of a model and a pilot study. *The Economic Journal* 127, F209–F235.
- Maritz, J. S. (2018). *Empirical Bayes methods with applications*. CRC Press.
- McLean, R. D. and J. Pontiff (2016). Does academic research destroy stock return predictability? *The Journal of Finance* 71(1), 5–32.
- Moskowitz, T. J., Y. H. Ooi, and L. H. Pedersen (2012). Time series momentum. *Journal of financial economics* 104(2), 228–250.
- Murtagh, F. and P. Legendre (2014). Ward’s hierarchical agglomerative clustering method: which algorithms implement ward’s criterion? *Journal of classification* 31(3), 274–295.
- Newey, W. K. and K. D. West (1987, May). A Simple, Positive Semi-definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix. *Econometrica* 55(3), 703–708.
- Nosek, B. A., J. R. Spies, and M. Motyl (2012). Scientific utopia: Ii. restructuring incentives and practices to promote truth over publishability. *Perspectives on Psychological Science* 7(6), 615–631.
- Pedersen, L. H., A. Babu, and A. Levine (2021). Enhanced portfolio optimization. *Financial Analysts Journal* 77(2), 124–151.
- Shumway, T. (1997). The delisting bias in crsp data. *The Journal of Finance* 52(1), 327–340.
- Sloan, R. G. (1996). Do stock prices fully reflect information in accruals and cash flows about future earnings? *Accounting Review* 71(3), 289–315.
- Ward, J. H. (1963). Hierarchical grouping to optimize an objective function. *Journal of the American Statistical Association* 58(301), 236–244.
- Welch, I. (2019). Reproducing, extending, updating, replicating, reexamining, and reconciling. *Critical Finance Review* 8(1-2), 301–304.

A Appendix: Additional Results and Proofs

Additional Results on Alpha Hacking

We consider the situation where the researcher has in-sample data from time 1 to time T and an out-of-sample (oos) period from time $T + 1$ to $T + T^{oos}$. The researcher may have used alpha-hacking during the in-sample period, but this does not affect the out-of-sample period. The researcher is interested in the posterior alpha based on the total evidence, in-sample and out-of-sample, which is useful for predicting factor performance in a future time period (that is, a time period that is out-of-sample relative to the existing out-of-sample period).

Proposition 6 (Out-of-sample alpha) *The posterior alpha based on an in-sample data from time 1 to T with alpha-hacking, and an out-of-sample period from $T + 1$ to $T + T^{oos}$ is given by*

$$E(\alpha|\hat{\alpha}, \hat{\alpha}^{oos}) = \kappa^{oos} (w(\hat{\alpha} - \bar{\varepsilon}) + (1 - w)\alpha^{oos}) \quad (\text{A.1})$$

where $w = \frac{\sigma^2/T^{oos}}{\bar{\sigma}^2/T + \sigma^2/T^{oos}} \in (0, 1)$ is the relative weight on the in-sample period relative to the out-of-sample period, and $\kappa^{oos} = \frac{1}{1 + 1/(\tau^2([\bar{\sigma}^2/T]^{-1} + [\sigma^2/T^{oos}]^{-1}))}$ is a shrinkage parameter.

We see that, the more alpha hacking the researcher has done (higher $\bar{\sigma}$), the less weight we put on the in-sample period relative to the out-of-sample period. Further, the in-sample period has the non-proportional discounting due to alpha hacking ($\bar{\varepsilon}$), which we don't have for out-of-sample evidence.

So this result formalizes the idea that an in-sample backtest plus live performance is *not* the same as a longer backtest. For example, 10 years of backtest plus 10 years of live performance is more meaningful than 20 years of backtest with no live performance. The difference is that the oos-performance is free from alpha-hacking.

Proofs and Lemmas

The proofs make repeated use of the following well-known property of multivariate Normally distributed random variable. If x and y are multivariate Normal:

$$\begin{bmatrix} x \\ y \end{bmatrix} \sim N \left(\begin{bmatrix} \mu_x \\ \mu_y \end{bmatrix}, \begin{bmatrix} \Sigma_{xx} & \Sigma_{yx} \\ \Sigma_{xy} & \Sigma_{yy} \end{bmatrix} \right) \quad (\text{A.2})$$

then the conditional distribution of x given y has the following Normal distribution:

$$x|y \sim N \left(\mu_x + \Sigma_{xy}\Sigma_{yy}^{-1}(y - \mu_y), \Sigma_{xx} - \Sigma_{xy}\Sigma_{yy}^{-1}\Sigma_{yx} \right) \quad (\text{A.3})$$

The proofs also make use of the following two lemmas.

Lemma 1 *For random variables x, y, z , it holds that $E(\text{Var}(x|y, z)) \leq E(\text{Var}(x|y))$ and, if the random variables are jointly normal, then $\text{Var}(x|y, z) \leq \text{Var}(x|y)$.*

Lemma 2 Let A be an $N \times N$ matrix for which all diagonal elements equal a and all off-diagonal elements equal b , where $a \neq b$ and $a + b(N - 1) \neq 0$. Then the inverse A^{-1} exists and is of the same form:

$$A = \begin{bmatrix} a & & b \\ & \ddots & \\ b & & a \end{bmatrix} \quad A^{-1} = \begin{bmatrix} c & & d \\ & \ddots & \\ d & & c \end{bmatrix} \quad (\text{A.4})$$

where $c = \frac{a+b(N-2)}{(a-b)(a+b(N-1))}$ and $d = \frac{-b}{(a-b)(a+b(N-1))}$.

Proof of Lemma 1. Using the definition of conditional variance, we have

$$E(\text{Var}(x|y, z)) = E(E(x^2|y, z)) - E([E(x|y, z)]^2) = E(x^2) - E([E(x|y, z)]^2)$$

Hence, using Jensen's inequality, we have

$$\begin{aligned} E(\text{Var}(x|y)) - E(\text{Var}(x|y, z)) &= E([E(x|y, z)]^2) - E([E(x|y)]^2) \\ &= E([E(x|y, z)]^2) - E([E(E(x|y, z)|y)]^2) \\ &\geq E([E(x|y, z)]^2) - E(E([E(x|y, z)]^2 | y)) = 0 \end{aligned}$$

The result for normal distributions follows from the fact that normal conditional variances are non-stochastic, i.e., $\text{Var}(x|y) = E(\text{Var}(x|y))$. In this case, we can also characterize the extra drop in variance due to conditioning on z using its orthogonal component ε from the regression $z = a + by + \varepsilon$, using similar notation as (A.2):

$$\begin{aligned} \text{Var}(x|y, z) = \text{Var}(x|y, \varepsilon) &= \Sigma_{x,x} - \Sigma_{x,(y,\varepsilon)} \Sigma_{(y,\varepsilon),(y,\varepsilon)}^{-1} \Sigma_{(y,\varepsilon),x} \\ &= \Sigma_{x,x} - \Sigma_{x,y} \Sigma_{y,y}^{-1} \Sigma_{y,x} - \Sigma_{x,\varepsilon} \Sigma_{\varepsilon,\varepsilon}^{-1} \Sigma_{\varepsilon,x} = \text{Var}(x|y) - \Sigma_{x,\varepsilon} \Sigma_{\varepsilon,\varepsilon}^{-1} \Sigma_{\varepsilon,x} \end{aligned}$$

■

Proof of Lemma 2. The proof follows from inspection: The product of A and its proposed inverse clearly has the same form as A with diagonal elements

$$ac + bd(I - 1) = \frac{a(a + b(N - 2)) - b^2(N - 1)}{(a - b)(a + b(N - 1))} = \frac{a^2 + ab(N - 1) - ab - b^2(N - 1)}{(a - b)(a + b(N - 1))} = 1$$

and off-diagonal elements

$$ad + bc + bd(N - 2) = \frac{-ab + b(a + b(N - 2)) - b^2(N - 2)}{(a - b)^2(a + b(N - 1))^2} = 0$$

In other words, AA^{-1} equals the identity, proving the result. ■

Proof of Equations (4)–(6). The posterior distribution of the true alpha given the observed

factor return is computed using (A.3). The conditional mean is

$$E(\alpha|\hat{\alpha}) = 0 + \frac{\text{Cov}(\alpha, \hat{\alpha})}{\text{Var}(\hat{\alpha})}(\hat{\alpha} - 0) = \frac{\tau^2}{\tau^2 + \sigma^2/T} \hat{\alpha} = \kappa \hat{\alpha}$$

where κ is given by (5) and the posterior variance is

$$\text{Var}(\alpha|\hat{\alpha}) = \text{Var}(\alpha) - \frac{(\text{Cov}(\alpha, \hat{\alpha}))^2}{\text{Var}(\hat{\alpha})} = \tau^2 - \tau^2 \frac{\tau^2}{\tau^2 + \sigma^2/T} = \frac{\tau^2 \sigma^2/T}{\tau^2 + \sigma^2/T} = \kappa \frac{\sigma^2}{T}$$

■

Proof of Proposition 1. The posterior alpha with alpha-hacking is given via (A.3) as

$$E(\alpha|\hat{\alpha}) = 0 + \frac{\text{Cov}(\alpha, \hat{\alpha})}{\text{Var}(\hat{\alpha})}(\hat{\alpha} - E(\hat{\alpha})) = \frac{\tau^2}{\tau^2 + \bar{\sigma}^2/T}(\hat{\alpha} - \bar{\varepsilon}) = -\kappa_0 + \kappa^{\text{hacking}} \hat{\alpha}$$

where $\kappa^{\text{hacking}} = \frac{1}{1 + \frac{\bar{\sigma}^2}{\tau^2 T}}$, $\kappa_0 = \kappa^{\text{hacking}} \bar{\varepsilon} \geq 0$, and $\kappa^{\text{hacking}} \leq \kappa$ because $\bar{\sigma} \geq \sigma$.

■

Proof of Proposition 2. The posterior mean given $\hat{\alpha}$ and $\hat{\alpha}^g$ is computed via (A.3) as

$$\begin{aligned} E(\alpha|\hat{\alpha}, \hat{\alpha}^g) &= [\tau^2 \quad \tau^2] \begin{bmatrix} \tau^2 + \sigma_T^2 & \tau^2 + \rho\sigma_T^2 \\ \tau^2 + \rho\sigma_T^2 & \tau^2 + \sigma_T^2 \end{bmatrix}^{-1} \begin{bmatrix} \hat{\alpha} \\ \hat{\alpha}^g \end{bmatrix} \\ &= \frac{1}{\det} [\tau^2 \quad \tau^2] \begin{bmatrix} \tau^2 + \sigma_T^2 & -(\tau^2 + \rho\sigma_T^2) \\ -(\tau^2 + \rho\sigma_T^2) & \tau^2 + \sigma_T^2 \end{bmatrix} \begin{bmatrix} \hat{\alpha} \\ \hat{\alpha}^g \end{bmatrix} \\ &= \frac{\tau^2(1 - \rho)\sigma_T^2}{\det} (\hat{\alpha} + \hat{\alpha}^g) \\ &= \frac{\tau^2(1 - \rho)}{\sigma_T^2(1 - \rho)(1 + \rho) + 2\tau^2(1 - \rho)} (\hat{\alpha} + \hat{\alpha}^g) \\ &= \kappa^g \left(\frac{1}{2} \hat{\alpha} + \frac{1}{2} \hat{\alpha}^g \right) \end{aligned}$$

using the notation $\sigma_T^2 = \sigma^2/T$ and

$$\det = (\tau^2 + \sigma_T^2)^2 - (\tau^2 + \rho\sigma_T^2)^2 = \sigma_T^2[\sigma_T^2(1 - \rho^2) + 2\tau^2(1 - \rho)].$$

The global shrinkage parameter κ^g is in $[\kappa, 1]$ and decreases with the correlation ρ , attaining the minimum value, $\kappa^g = \kappa$, when $\rho = 1$ as is clearly seen from (12).

The result about the posterior variance follows from Lemma 1.

■

Proof of Proposition 3. The prior joint distribution of the true and estimated alphas is

given by the following expression, where we focus on factor 1 without loss of generality:

$$\begin{bmatrix} \alpha^1 \\ \hat{\alpha}^1 \\ \vdots \\ \hat{\alpha}^N \end{bmatrix} \sim N \left(\begin{bmatrix} 0 \\ 0 \\ \vdots \\ 0 \end{bmatrix}, \begin{bmatrix} \tau_c^2 + \tau_w^2 & \tau_c^2 + \tau_w^2 & \tau_c^2 & \cdots & \tau_c^2 \\ \tau_c^2 + \tau_w^2 & \tau_c^2 + \tau_w^2 + \sigma^2/T & & & \tau_c^2 + \rho\sigma^2/T \\ \tau_c^2 & & & & \\ \vdots & & & \ddots & \\ \tau_c^2 & \tau_c^2 + \rho\sigma^2/T & & & \tau_c^2 + \tau_w^2 + \sigma^2/T \end{bmatrix} \right) \quad (\text{A.5})$$

The posterior alpha of factor 1 is therefore normally distributed with a mean derived using the standard formula for conditional normal distributions (A.3):

$$E(\alpha^1 | \hat{\alpha}^1, \dots, \hat{\alpha}^N) = \begin{bmatrix} \tau_c^2 + \tau_w^2 \\ \tau_c^2 \\ \vdots \\ \tau_c^2 \end{bmatrix}^\top \begin{bmatrix} \tau_c^2 + \tau_w^2 + \sigma^2/T & & & \\ & \tau_c^2 + \rho\sigma^2/T & & \\ & & \ddots & \\ \tau_c^2 + \rho\sigma^2/T & & & \tau_c^2 + \tau_w^2 + \sigma^2/T \end{bmatrix}^{-1} \begin{bmatrix} \hat{\alpha}^1 \\ \vdots \\ \hat{\alpha}^N \end{bmatrix}$$

We next use Lemma 2 and its notation, i.e., $a = \tau_c^2 + \tau_w^2 + \sigma^2/T$, $b = \tau_c^2 + \rho\sigma^2/T$, and c', d are defined accordingly, where we use the notation c' to avoid confusion with the c in equation (14). This application of Lemma 2 yields

$$\begin{aligned} E(\alpha^1 | \hat{\alpha}^1, \dots, \hat{\alpha}^N) &= \begin{bmatrix} \tau_c^2 + \tau_w^2 \\ \tau_c^2 \\ \vdots \\ \tau_c^2 \end{bmatrix}^\top \begin{bmatrix} c' & d \\ & \ddots \\ d & c' \end{bmatrix} \begin{bmatrix} \hat{\alpha}^1 \\ \vdots \\ \hat{\alpha}^N \end{bmatrix} \\ &= \begin{bmatrix} \tau_c^2(c' + d(N-1)) + \tau_w^2 c' \\ \tau_c^2(c' + d(N-1)) + \tau_w^2 d \\ \vdots \\ \tau_c^2(c' + d(N-1)) + \tau_w^2 d \end{bmatrix}^\top \begin{bmatrix} \hat{\alpha}^1 \\ \vdots \\ \hat{\alpha}^N \end{bmatrix} \\ &= (\tau_c^2(c' + d(N-1)) + \tau_w^2 d) N \hat{\alpha}^\cdot + \tau_w^2(c' - d) \hat{\alpha}^1 \\ &= \left(\tau_c^2 \frac{N}{a + b(N-1)} - \tau_w^2 \frac{bN}{(a-b)(a+b(N-1))} \right) \hat{\alpha}^\cdot + \tau_w^2 \frac{1}{a-b} \hat{\alpha}^1 \\ &= \frac{\tau_c^2}{b + \frac{a-b}{N}} \hat{\alpha}^\cdot + \frac{\tau_w^2}{a-b} \left(\hat{\alpha}^1 - \frac{1}{1 + \frac{a-b}{bN}} \hat{\alpha}^\cdot \right) \\ &= \frac{\tau_c^2}{\tau_c^2 + \rho\sigma^2/T + \frac{\tau_w^2 + (1-\rho)\sigma^2/T}{N}} \hat{\alpha}^\cdot + \frac{\tau_w^2}{\tau_w^2 + (1-\rho)\sigma^2/T} \left(\hat{\alpha}^1 - \frac{1}{1 + \frac{\tau_w^2 + (1-\rho)\sigma^2/T}{(\tau_c^2 + \rho\sigma^2/T)N}} \hat{\alpha}^\cdot \right) \\ &= \frac{1}{1 + \frac{\rho\sigma^2}{\tau_c^2 T} + \frac{\tau_w^2 + (1-\rho)\sigma^2/T}{\tau_c^2 N}} \hat{\alpha}^\cdot + \frac{1}{1 + \frac{(1-\rho)\sigma^2}{\tau_w^2 T}} \left(\hat{\alpha}^1 - \frac{1}{1 + \frac{\tau_w^2 + (1-\rho)\sigma^2/T}{(\tau_c^2 + \rho\sigma^2/T)N}} \hat{\alpha}^\cdot \right) \end{aligned}$$

The posterior has conditional variance

$$\begin{aligned}
\text{Var}(\alpha^1 | \hat{\alpha}^1, \dots, \hat{\alpha}^N) &= \tau_c^2 + \tau_w^2 - \begin{bmatrix} \tau_c^2 + \tau_w^2 \\ \tau_c^2 \\ \vdots \\ \tau_c^2 \end{bmatrix}^\top \begin{bmatrix} c' & d \\ & \ddots \\ d & c' \end{bmatrix} \begin{bmatrix} \tau_c^2 + \tau_w^2 \\ \tau_c^2 \\ \vdots \\ \tau_c^2 \end{bmatrix} \\
&= \tau_c^2 + \tau_w^2 - \begin{bmatrix} \tau_c^2(c' + d(N-1)) + \tau_w^2 c' \\ \tau_c^2(c' + d(N-1)) + \tau_w^2 d \\ \vdots \\ \tau_c^2(c' + d(N-1)) + \tau_w^2 d \end{bmatrix}^\top \begin{bmatrix} \tau_c^2 + \tau_w^2 \\ \tau_c^2 \\ \vdots \\ \tau_c^2 \end{bmatrix} \\
&= \tau_c^2 + \tau_w^2 - (\tau_c^2(c' + d(N-1)) + \tau_w^2 c')(\tau_c^2 + \tau_w^2) \\
&\quad - (\tau_c^2(c' + d(N-1)) + \tau_w^2 d)\tau_c^2(N-1) \\
&\rightarrow \tau_c^2 + \tau_w^2 - \left(\tau_c^2 \left(\frac{1}{a-b} - \frac{1}{a-b} \right) + \tau_w^2 \frac{1}{a-b} \right) (\tau_c^2 + \tau_w^2) \\
&\quad - \left(\tau_c^2 \frac{1}{b} - \tau_w^2 \frac{1}{a-b} \right) \tau_c^2 \\
&= \tau_c^2 + \tau_w^2 - \left(\tau_w^4 \frac{1}{a-b} + \tau_c^4 \frac{1}{b} \right) \\
&= \tau_c^2 + \tau_w^2 - \left(\frac{\tau_w^4}{\tau_w^2 + (1-\rho)\sigma^2/T} + \frac{\tau_c^4}{\tau_c^2 + \rho\sigma^2/T} \right)
\end{aligned}$$

The last results follow from Lemma 1. ■

Proof of Proposition 4. We write the joint prior distribution of true and observed alphas in the multi-level hierarchical model as

$$\begin{pmatrix} \alpha \\ \hat{\alpha} \end{pmatrix} \sim N \left(\alpha^0 \mathbf{1}_{2NK}, \begin{pmatrix} \Omega & \Omega \\ \Omega & \Omega + \Sigma/T \end{pmatrix} \right) \quad (\text{A.6})$$

The posterior mean vector of true alphas is computed via (A.3):

$$\begin{aligned}
E(\alpha | \hat{\alpha}) &= \mathbf{1}_{NK} \alpha_0 + \Omega (\Omega + \Sigma/T)^{-1} (\hat{\alpha} - \mathbf{1}_{NK} \alpha_0) \\
&= (\Omega^{-1} + T\Sigma^{-1})^{-1} (\Omega^{-1} \mathbf{1}_{NK} \alpha_0 + T\Sigma^{-1} \hat{\alpha})
\end{aligned}$$

using that $(\Omega + \Sigma/T)^{-1} = \Omega^{-1} - \Omega^{-1} (\Omega^{-1} + T\Sigma^{-1})^{-1} \Omega^{-1}$ by the Woodbury matrix identity. The posterior variance is computed similarly via (A.3) and the same application of the Woodbury matrix identity as

$$\text{Var}(\alpha | \hat{\alpha}) = \Omega - \Omega (\Omega + \Sigma/T)^{-1} \Omega = (\Omega^{-1} + T\Sigma^{-1})^{-1}.$$
■

Proof of Proposition 5. Based on the definition of the Bayesian FDR, we have:

$$\begin{aligned}
\text{FDR}^{\text{Bayes}} &= E \left(\frac{\sum_i 1_{\{i \text{ false discovery}\}}}{\sum_i 1_{\{i \text{ discovery}\}}} \middle| \hat{\alpha}^1, \dots, \hat{\alpha}^N, \tau \right) \\
&= \frac{1}{\sum_i 1_{\{i \text{ discovery}\}}} E \left(\sum_i 1_{\{i \text{ false discovery}\}} \middle| \hat{\alpha}^1, \dots, \hat{\alpha}^N, \tau \right) \\
&= \frac{1}{\sum_i 1_{\{i \text{ discovery}\}}} \sum_i \text{Pr}(i \text{ false discovery} | \hat{\alpha}^1, \dots, \hat{\alpha}^N, \tau) \\
&= \frac{1}{\#\text{discoveries}} \sum_{i \text{ discovery}} p\text{-null}_i \\
&\leq 2.5\%
\end{aligned} \tag{A.7}$$

Proof of Proposition 6. The posterior mean alpha is

$$\begin{aligned}
E(\alpha | \hat{\alpha}, \hat{\alpha}^{\text{oos}}) &= [\tau^2 \quad \tau^2] \begin{bmatrix} \tau^2 + \bar{\sigma}_T^2 & \tau^2 \\ \tau^2 & \tau^2 + \sigma_{\text{oos}}^2 \end{bmatrix}^{-1} \begin{bmatrix} \hat{\alpha} - \bar{\varepsilon} \\ \hat{\alpha}^{\text{oos}} \end{bmatrix} \\
&= \frac{1}{\det} [\tau^2 \quad \tau^2] \begin{bmatrix} \tau^2 + \sigma_{\text{oos}}^2 & -\tau^2 \\ -\tau^2 & \tau^2 + \bar{\sigma}_T^2 \end{bmatrix} \begin{bmatrix} \hat{\alpha} - \bar{\varepsilon} \\ \hat{\alpha}^{\text{oos}} \end{bmatrix} \\
&= \frac{\tau^2}{\det} (\sigma_{\text{oos}}^2 (\hat{\alpha} - \bar{\varepsilon}) + \bar{\sigma}_T^2 \hat{\alpha}^{\text{oos}}) \\
&= \frac{\tau^2 (\bar{\sigma}_T^2 + \sigma_{\text{oos}}^2)}{\tau^2 (\bar{\sigma}_T^2 + \sigma_{\text{oos}}^2) + \bar{\sigma}_T^2 \sigma_{\text{oos}}^2} (w(\hat{\alpha} - \bar{\varepsilon}) + (1 - w)\alpha^{\text{oos}}) \\
&= \frac{\tau^2}{\tau^2 + \bar{\sigma}_T^2 \sigma_{\text{oos}}^2 / (\bar{\sigma}_T^2 + \sigma_{\text{oos}}^2)} (w(\hat{\alpha} - \bar{\varepsilon}) + (1 - w)\alpha^{\text{oos}}) \\
&= \frac{1}{1 + \frac{1}{\tau^2 (\bar{\sigma}_T^2 + \sigma_{\text{oos}}^2)}} (w(\hat{\alpha} - \bar{\varepsilon}) + (1 - w)\alpha^{\text{oos}})
\end{aligned}$$

using the notation $\bar{\sigma}_T^2 = \bar{\sigma}^2/T$, $\sigma_{\text{oos}}^2 = \sigma^2/T^{\text{oos}}$, and

$$\det = (\tau^2 + \bar{\sigma}_T^2)(\tau^2 + \sigma_{\text{oos}}^2) - \tau^4 = \tau^2(\bar{\sigma}_T^2 + \sigma_{\text{oos}}^2) + \bar{\sigma}_T^2 \sigma_{\text{oos}}^2.$$

B Empirical Bayes Estimation

For convenient reference, we restate the multi-level hierarchical model of Section 1. For a factor i in cluster j and corresponding to signal n , the factor is

$$f_t^i = \alpha^i + \beta^i r_t^m + \varepsilon_t^i$$

with

$$\alpha^i = \alpha^o + c^j + s^n + w^i$$

where the alpha components are $\alpha^o = 0$, $c^j \sim N(0, \tau_c^2)$, $s^n \sim N(0, \tau_s^2)$, and $w^i \sim N(0, \tau_w^2)$. We write alpha in vector form as

$$\alpha = \alpha^o \mathbf{1}_{NK} + Mc + Zs + w \tag{B.1}$$

where $\alpha = (\alpha^1, \dots, \alpha^{NK})'$, $c = (c^1, \dots, c^J)'$, $s = (s^1, \dots, s^N)'$, $w = (w^1, \dots, w^{NK})'$, M is the $NK \times J$ matrix of cluster memberships, and Z is the $NK \times N$ matrix indicating the characteristic that factor i is based on. Given the hyperparameters $(\alpha^o, \tau_c, \tau_s, \tau_w)$, the prior mean and covariance matrix of alphas are

$$E[\alpha] = 0, \quad \Omega \equiv \text{Var}(\alpha) = MM'\tau_c^2 + ZZ'\tau_s^2 + I_{NK}\tau_w^2. \tag{B.2}$$

The vector of return shocks is $\varepsilon_t = (\varepsilon_t^1, \dots, \varepsilon_t^{NK})'$ which is distributed $\varepsilon_t \sim N(0, \Sigma)$.

Given this structure, we estimate the model as follows. The vector of factor returns $f_t = (f_t^1, \dots, f_t^{NK})'$ has marginal likelihood—that is, after integrating out the uncertain alpha components—that is distributed as

$$f_t \sim N(0, [\Omega + \Sigma])$$

or, equivalently (treating CAPM betas as known), the estimated alphas are distributed⁴⁷

$$\hat{\alpha} \sim N(0, [\Omega + \Sigma/T]).$$

The matrices Z and M are given by the factor definition and cluster assignment (Table J.2), respectively. We use a plug-in estimate of the factor CAPM-residual return covariance matrix, denoted $\hat{\Sigma}$ (discussed below). Finally, given $\hat{\Sigma}$, Z , and M , we estimate the hyperparameters of the prior distribution, (τ_c, τ_s, τ_w) via MLE based on the marginal likelihood.

This estimation approach is an example of the empirical Bayes method. It approximates the fully Bayesian posterior calculation (which requires integrating over a hyperprior distribution of hyperparameters, usually an onerous calculation) by setting the hyperparameters to their most likely values based on the marginal likelihood. It is particularly well suited to hierarchical Bayesian models in which parameters for individual observations share some common structure, so that the realized heterogeneity across individual is informative about sensible values for the hyperparameters of the prior. Our model and estimation approach implementation is a minor variation on Bayesian hierarchical normal mean models that are common in Bayesian statistics (textbook treatments include Efron, 2012; Gelman et al., 2013; Maritz, 2018). We conduct sensitivity analysis to ensure that our results are robust to a wide range of hyperparameters (see Figure F.1). Also, we note that our EB methodology is more easily replicable than a full-Bayesian setting with additional hyperpriors as EB relies on a closed-form Bayesian updating rather than a numerical integration.

⁴⁷We abstract from uncertainty in CAPM betas to emphasize the Bayesian updating of alphas. Our conclusions are qualitatively insensitive to accounting for beta uncertainty.

To ensure cross-sectional stationarity, we scale each factor such that their monthly idiosyncratic volatility is $10\%/\sqrt{12}$ (i.e., 10% annualized). To construct a plug-in estimate of the factor residual return covariance matrix, denoted $\hat{\Sigma}$, we face two main empirical challenges. First, the sample covariance is poorly behaved due the relatively large number of factors compared to the number of time series observations. Second, we have an unbalanced panel because different factors come online at different points in time. To address the first challenge, we impose a block equicorrelation structure on Σ based on factors' cluster membership.⁴⁸ The correlation between factors in clusters i and j is estimated as the average correlation among all pairs such that one factor is in cluster i and the other is in j . In our global analyses, blocks correspond to region-cluster pairs. To address unbalancedness, we use the bootstrap. In particular, we generate 10,000 bootstrap samples that resample rows of the unbalanced factor return dataset. Each bootstrap sample is, therefore, also unbalanced, and we use this to produce a distribution of alpha estimates. From this we calculate $\hat{\Sigma}/T$ as the covariance of alphas across bootstrap samples (imposing the block equicorrelation structure).

Table B.1 shows the estimated hyperparameters across different samples. While most of our analysis is based on these full-sample estimates, we also consider rolling-estimates of when considering out-of-sample evidence as seen in Figure D.4.

Sample	τ_c	τ_w	τ_s
USA	0.35%	0.21%	
Developed	0.24%	0.18%	
Emerging	0.32%	0.24%	
USA, Developed & Emerging	0.30%	0.19%	0.10%
World	0.37%	0.23%	
World ex. US	0.29%	0.20%	
USA - Mega	0.26%	0.16%	
USA - Large	0.31%	0.18%	
USA - Small	0.44%	0.26%	
USA - Micro	0.48%	0.32%	
USA - Nano	0.42%	0.28%	

Table B.1: Hyperparameters of the prior distribution estimated by maximum likelihood. Here, τ_c is the estimated dispersion in cluster alphas (e.g., the dispersion in the alpha of the value cluster alpha, momentum cluster, and so on). When we estimate a single region, τ_w is the idiosyncratic dispersion of alphas within each cluster. When we jointly estimate several regions, then τ_s is the estimated dispersion in alphas across signals within each cluster, and τ_w is the estimated idiosyncratic dispersion in alphas for factors identified by their signal and region.

⁴⁸As advocated by Engle and Kelly (2012) and Elton et al. (2006), block equicorrelation constrains all pairs of factors in the same block to share a single correlation parameter, and likewise for cross-block correlations. This stabilizes covariance matrix estimates by dramatically reducing the parameterization of the correlation matrix, while leaving the individual variance estimates unchanged.

Internet Appendix

C Differences in Sample and Factor Construction

Here we provide further details on the difference in sample and factor construction vs. [Hou et al. \(2020\)](#) accounting for the difference between the first two bars in Figure 1 in our introduction. To re-iterate, with raw returns and capped value weights, we find a replication rate of 55.6% where as [Hou et al. \(2020\)](#) finds a replication rate of 35%.

This difference has the following decomposition.⁴⁹ First, [Hou et al. \(2020\)](#) focus their analysis on value-weighted factors rather than the standard [Fama and French \(1993\)](#) methodology that gives half the weight to small stocks (or equal-weighting that gives even more weight to small stocks). However, pure value weighting sometimes leads to excessively concentrated portfolios that mask the behavior of factors.⁵⁰ We use a weighting scheme that we refer to as “capped value-weighting” that winsorizes market caps at the NYSE 80th percentile. This weighting is a helpful compromise between pure value-weighting and the Fama-French method since our factors continue to emphasize large stocks, but the capped scheme avoids undue skewness toward a few mega stocks, which in turn produces more robust factor behavior over time and across countries. Capped value weights contribute +9.2% to our higher replication rate. Figure C.1 reproduces Figure 1 with straight value weights.

Second, for each characteristic, [Hou et al. \(2020\)](#) construct three variations on each factor having either 1-month, 6-month, or 12-month holding periods. They treat these as separate factors so that their factor count essentially multiplies their characteristics count by a factor of three. In contrast, we focus on 1-month returns because this is the horizon of interest in almost all of the original studies (and we believe it is the most economically meaningful since it uses the most current data as theory dictates). Our focus on only the 1-month holding period factor for each characteristic contributes +5.0% to our replication rate.

Third, we use a longer sample, which contributes +8.3% to the difference in replication rate. Fourth, we add 15 factors to our sample that were previously studied in the literature but not studied by [Hou et al. \(2020\)](#), which has a no effect on the replication rate.

Finally, we use tercile spreads and breakpoints based on all stocks above the NYSE 20th percentile (i.e., non-micro-caps), while [Hou et al. \(2020\)](#) use decile spreads and breakpoints based on all NYSE stocks. Our more conservative method leads to a -6.0% drop in the replication rate. The remaining +4.1% difference in replication rates is due to minor construction and sample details⁵¹. We discuss this decomposition further in Section 2, where we detail our factor construction choices and discuss why we prefer them.

⁴⁹Note that the attribution to specific changes depends on the order in which the changes are applied.

⁵⁰For example, Nokia stock accounted for more than 70% of the total market capitalization in Finland in 1999 and 2000.

⁵¹For example we always lag accounting data four months, they use a mixture of updating schemes and our set of factors is not identical to that in [Hou et al. \(2020\)](#).

Replication Rate with Uncapped Value Weights

In Figure C.1, we show an alternative version of Figure 1 with factors constructed using straight (as opposed to capped) value weights. It shows that all of main our conclusions remain similar. Our ultimate replication rate in this case is 79.8% (based on global data and Bayesian model estimates).




Figure C.1: Replication Rates Versus the Literature (Uncapped Value-weighting)

Note: This figure reproduces the analysis of figure 1 using uncapped value weights to construct factors.

D Additional Time-Series Results




Figure D.1: Out-of-sample performance of significant factors under empirical Bayes

Note: The figure shows the cumulative CAPM alpha of an average of factors significant under our empirical Bayes framework. The significance cutoffs are re-estimated each year with the available data. Factors are eligible for inclusion after the sample period in the original paper, so all returns are out-of-sample. The table shows the information ratio (alpha divided by residual volatility) for the full sample (1990-2020) with t -statistics in parentheses.




Figure D.2: In-Sample versus Out-of-Sample Alphas for US Factors

Note: The figure plots OLS alphas for US factors during the in-sample period (i.e., the period studied in the original publication) versus out-of-sample alphas. In Panel A, out-of-sample is the time period before the in-sample period. In Panel B, out-of-sample is the time period before the in-sample period. In Panel C, out-of-sample includes both the time period before and after the in-sample period. We require at least five years of out-of-sample data for a factor to be included, amounting to 102, 115 and 119 factors in panel A, B and C. The figure also reports feasible GLS estimates of out-of-sample alphas on in-sample alphas and in-sample alphas squared. To implement feasible GLS, we assume that the error variance-covariance matrix is proportional to the full-sample CAPM residual variance-covariance matrix, $\hat{\Sigma}/T$, described in section B. The blue line is a local polynomial regression fit where observations are weighted by their vicinity to the point on the x-axis. The shaded area is 95% confidence bands. The dotted line is the 45° line.

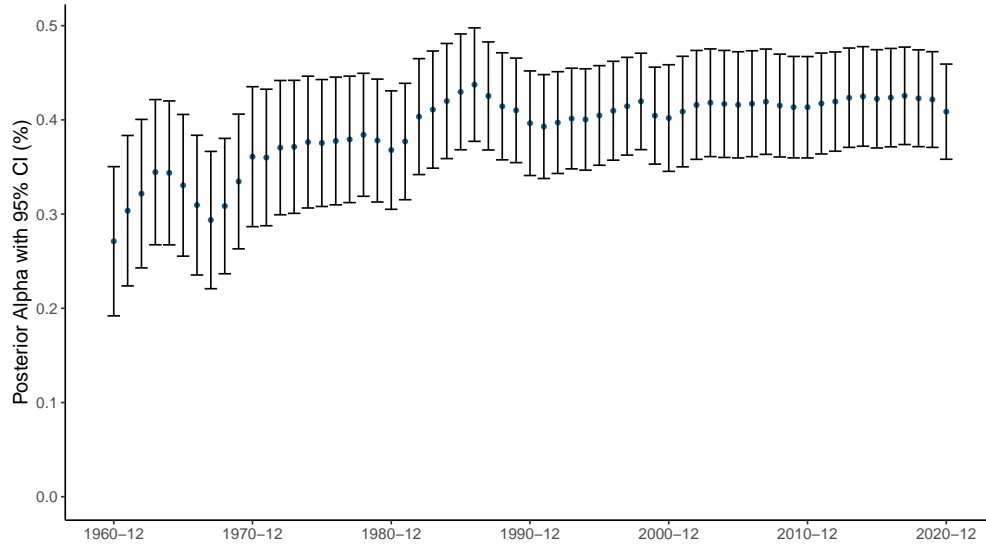


Figure D.3: World Factor Alpha Posterior Distribution Over Time

Note: The figure reports the CAPM alpha and 95% posterior confidence interval for an equal-weighted portfolio of World factors based on EB posteriors re-estimated in December each year. In contrast to figure 9, we re-estimate τ_c and τ_w at each point in time. Figure D.4 shows how the estimated taus evolve over time.




Figure D.4: World Factor Hyperparameters Over Time

Note: The figure reports the τ_c and τ_w used in figure D.3.

E Economic Benefit of More Powerful Tests

Table E.1: The Economic Benefit of More Powerful Tests

	Region		
	US	Developed ex. US	Emerging
	(1)	(2)	(3)
Alpha	0.35*** (5.05)	0.24*** (5.33)	0.27*** (3.66)
Market Beta	-0.12*** (-4.33)	-0.09*** (-5.68)	-0.04*** (-3.14)
Observations	540	420	388
Adjusted R^2	0.17	0.18	0.03

Note: The dependent variable is an equal-weighted portfolio of factors that are significant under empirical Bayes (EB), but not under OLS with the Benjamini-Yekutieli adjustment (BY). A factor is significant under EB when the probability of a negative alpha is below 2.5%. A factor is significant under BY when the adjusted two-sided p -value is below 5%, and the OLS alpha estimate is positive. Starting in 1959, we update the posterior distribution and the OLS estimates by the end of each year and invest in marginally significant factors over the subsequent year. To avoid lookahead bias, we only use factors after the sample in the original paper has ended. We only consider factors found to be significant by the original reference. The alpha estimates are in percentages, with t -statistics in parentheses. Standard errors are computed following Newey and West (1987) with 6 lags. The stars indicate * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

F Accounting for Publication Bias

Harvey et al. (2016) provides a framework to estimate the total number of factors researchers have tried. The framework is based on t -statistics of published factors and estimation framework to determine the number of unobserved factors.

One set of simulations is constructed to match the baseline scenario of (Harvey et al., 2016, Table 5.A, row 1), which estimates that researchers have tried $M = 1,297$ factors, of which 39.6% of have zero alpha and another is based on the more conservative scenario of (Harvey et al., 2016, Table 5.B, row 1), which implies that researchers have tried 2458 factors, of which 68.3% have zero alpha. Harvey et al. (2016) states that “the average annual Sharpe ratio for these [true] factors is 0.44.”

To incorporate these unobserved factors into our framework, we proceed as follows for the baseline scenario. We simulate a total of 1,300 factors in 26 clusters of 50 factors per cluster. We let all factors in 10 clusters have true alphas equal to zero while the remaining clusters have non-zero true alphas. For each of the clusters with non-zero alphas, we set the cluster alpha to $c^j = 0.44 \times 10\% / 12$ so that the monthly abnormal return corresponds to an annual Sharpe ratio of 0.44 given the annual volatility of 10%. Finally, we draw each factor’s true alpha from $\alpha^i \sim N(c^j, \tau_w^2)$, and then simulate 70 years of monthly returns with within-cluster correlation of 0.58 and 0.02 otherwise.⁵² Finally, we estimate prior parameters τ using this data with the same method that we used on the observed data. We repeat this simulation process and compute the average τ_c , which is interpreted as a value that accounts for unobserved factors of the form implied by Harvey et al. (2016). We note that we are implicitly assuming that the unobserved factors belong to different clusters, such that observing new poor performing factors would lead to more shrinkage toward zero via a lower τ_c , but not via different cluster mean returns.

Similarly for the conservative scenario, we simulate a total of 2500 factors in 50 clusters of 50 factors per cluster. We let all factors in 16 clusters have true alphas equal to zero while the remaining clusters have non-zero true alphas as described above. Figure F.1 shows the US replication rate under these alternative hyperparameters of the prior distribution.

⁵²The values are calibrated to match the data on US factors.




Figure F.1: Replication Rate with Prior Estimated in Light of Unobserved Factors

Note: The figure shows how the replication rate in the US varies when changing the τ_c and τ_w parameter. The dotted line shows our replication rate of 82.4%. The data estimate of τ_w is 0.21%. The green square, highlights the value estimated in the data $\tau_c = 0.35\%$. The red triangle and the blue circle highlights values that are found by estimating the empirical Bayes model according to assumptions about unobserved factors from Harvey et al. (2016). The values are $\tau_c = 0.28\%$ in the baseline scenario and $\tau_c = 0.20\%$ in the conservative scenario. A description of this approach can be found in the appendix, section F.

G Results by Cluster, Region, and Size

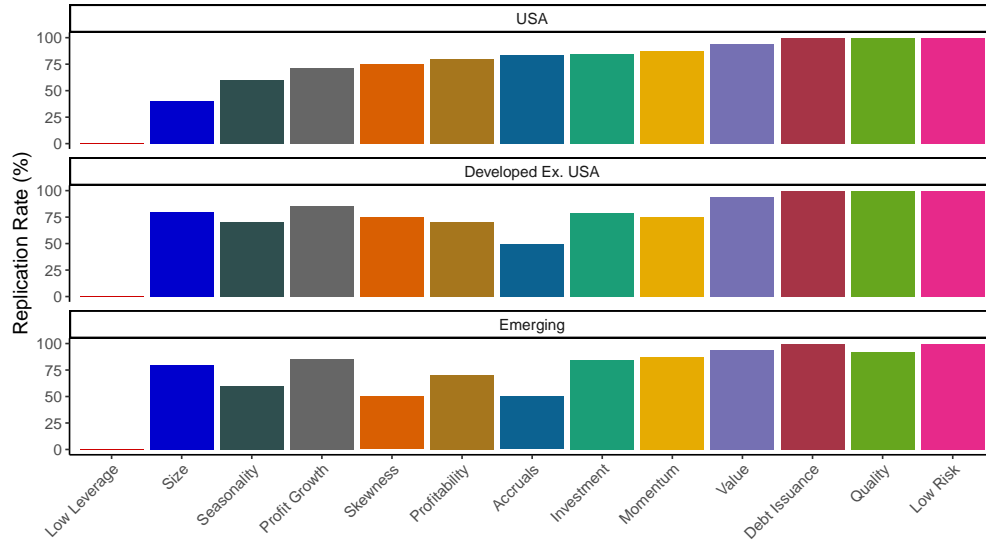


Figure G.1: Replication Rates across Regions by Cluster

Note: Share of factors within each cluster where the 95% posterior intervals does not include zero.




Figure G.2: Replication Rates across Size Groups by Cluster

Note: The figure shows replication rates for US factors created within a size group using rank weights. Mega stocks have a market cap higher than the 80th percentile of NYSE stocks, large stocks are between the 80th and 50th percentile, small stocks are between the 50th and 20th percentile, micro stocks are between the 20th and 1st percentile and nano stocks have a market cap below the 1st percentile of NYSE stocks.

H Further Results on the Tangency Portfolio

In this section, we elaborate on the influence of factors on the tangency portfolio (TPF). Figure H.1 and H.2 shows the tangency weights across regions and size groups. Most notably, the size cluster is much more important outside of the US and among smaller stocks.

The previous TPF analysis, and the results shown in figure 13 has used the 13 cluster level portfolios in addition to the market portfolio as inputs. In the remainder of this section, we build the TPF at the factor level by using the 119 US factors that were found to be significant by an earlier paper. We start the analysis in 1972 to ensure that all factors have non-missing data. The main issue with a factor level analysis, is estimating the covariance matrix. We follow Pedersen et al. (2021) and adjust the covariance matrix by shrinking the correlations towards zero

$$\Sigma_w = \sigma[(1 - w)\mathbf{\Omega} + w\mathbf{I}]\sigma$$

where $\mathbf{\Omega}$ is the sample correlation matrix, σ is a matrix with the sample volatilities on the diagonal and zero elsewhere, \mathbf{I} is the identity matrix and w is a shrinkage parameter. The tangency weights are recovered from the standard formula on the adjusted covariance matrix. This approach requires choosing the shrinkage parameter. Ideally, we want to choose the amount of shrinkage to maximize out-of-sample Sharpe ratio. We implement this intuition via. five-fold cross validation. In each fold, we estimate the tangency weights with a given shrinkage parameter using 4/5 of the data, and compute the realized Sharpe ratio on the remaining 1/5. We repeat this procedure 5 times, and compute the average realized Sharpe ratio for $w \in (0, 0.1, \dots, 1)$. In unreported results, we find that the optimal shrinkage parameter is $w = 0.5$ ⁵³.

Figure H.3 shows the in-sample Sharpe ratio of the tangency portfolio that are allowed to invest in the market portfolio and factors from one clusters. The dashed line shows the Sharpe ratio of the market portfolio. Figure H.4 shows the in-sample Sharpe ratio attainable after excluding factors from one cluster at a time. Figure H.5 shows the importance of each factor for the cluster TPF. Specifically, we report the drop in the maximal attainable Sharpe ratio within a cluster after excluding one of the cluster factors. Finally, figure H.6 shows the importance of each factor for the TPF that includes all factors. Specifically, we eliminate each factor one at a time and record the resulting drop in the in-sample Sharpe ratio.

⁵³The average monthly out-of-sample Sharpe ratio with $w = 0.5$ is 0.56 compared to 0.43 from the unconstrained solution ($w = 0$).




Figure H.1: Tangency Portfolio Weights across Regions

Note: Within each region, we compute the cluster return as the equal weighted return of all factors with data available at a given point in time. We further add the regional market return. We estimate the tangency weights following the method of Britten-Jones (1999) with a non-negativity constraint. The error bars are the 90% confidence intervals based on 10,000 bootstrap samples and the percentile method. The data starts in 1952 for the US, 1987 for Developed ex. US and 1994 for Emerging.




Figure H.2: Tangency Portfolio Weights across Size Groups

Note: Within each size group, we compute the cluster return as the equal weighted return of all factors with data available at a given point in time. We only use US data. We add the US market return. We estimate the tangency weights following the method of Britten-Jones (1999) with a non-negativity constraint. The error bars are the 90% confidence intervals based on 10,000 bootstrap samples and the percentile method. The data starts in 1963.




Figure H.3: Market + Cluster

Note: Each bar shows the monthly in-sample Sharpe ratio of a tangency portfolio that is allowed to invest in all factor from one cluster plus the market portfolio. We use the simple enhanced portfolio optimization method from Pedersen et al. (2021), with a shrinkage parameter of $w = 0.5$. The analysis is done on US factors from 1972 to 2020.




Figure H.4: Excluding One Cluster

Note: Each bar shows the monthly in-sample Sharpe ratio of a tangency portfolio that is allowed to invest in the market portfolio and factors from all clusters except one. We use the simple enhanced portfolio optimization method from Pedersen et al. (2021), with a shrinkage parameter of $w = 0.5$. The analysis is done on US factors from 1972 to 2020.

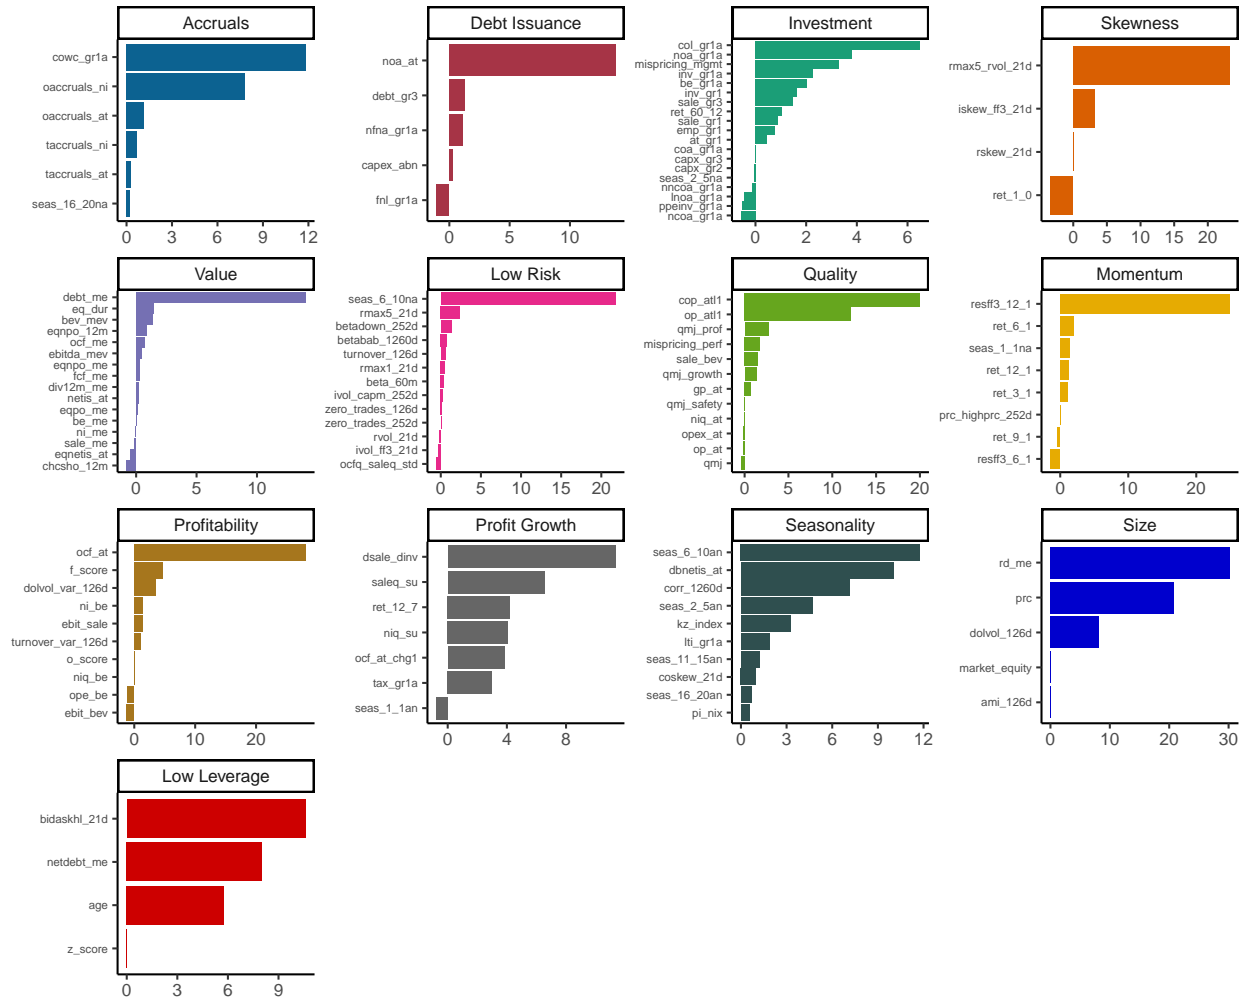


Figure H.5: Factor importance for cluster TPF

Note: Each bar shows the difference in the monthly in-sample Sharpe ratio of a tangency portfolio that invest in all factors within a cluster and a tangency portfolio that invest in all cluster factors except one. We show individual factors by their cluster. We use the simple enhanced portfolio optimization method from Pedersen et al. (2021), with a shrinkage parameter of $w = 0.5$. The analysis is done on US factors from 1972 to 2020.




Figure H.6: Factor importance for full TPF

Note: Each bar shows the difference in the monthly in-sample Sharpe ratio of a tangency portfolio that invest in all factors and a tangency portfolio that invest in all factor except one. We show individual factors by their cluster. We use the simple enhanced portfolio optimization method from Pedersen et al. (2021), with a shrinkage parameter of $w = 0.5$. The analysis is done on US factors from 1972 to 2020.

I Cluster Construction




Figure I.1: Clustering Factors into Themes

Note: This figure shows a hierarchical clustering of all factors into 13 themes using the sample of US stocks from 1975-2020. Long high indicates whether the factor is long stocks with a high value of the underlying characteristic.




Figure I.2: Factor Theme Correlations

Note: This figure shows the average pairwise Pearson correlation between factors from different clusters (off diagonal elements) or between factors in the same cluster (diagonal elements), using data on US stocks during the 1975-2020 period.

J Details on Clusters, Factors, and Countries

Table J.1: Factor and Cluster Details

Description	Variable Name	Citation	Orig. Sample	Sign	Orig. Signif.
Accruals					
Change in current operating working capital	cowc_gr1a	Richardson et al. (2005)	1962-2001	-1	1
Operating accruals	oaccruals_at	Sloan (1996)	1962-1991	-1	1
Percent operating accruals	oaccruals_ni	Hafzalla Lundholm and Van Winkle (2011)	1989-2008	-1	1
Years 16-20 lagged returns, nonannual	seas_16_20na	Heston and Sadka (2008)	1965-2002	-1	1

Total accruals	taccruals_at	Richardson et al. (2005)	1962-2001	-1	1
Percent total accruals	taccruals_ni	Hafzalla Lundholm and Van Winkle (2011)	1989-2008	-1	1

Debt Issuance

Abnormal corporate investment	capex_abn	Titman Wei and Xie (2004)	1973-1996	-1	1
Growth in book debt (3 years)	debt_gr3	Lyandres Sun and Zhang (2008)	1970-2005	-1	1
Change in financial liabilities	fnl_gr1a	Richardson et al. (2005)	1962-2001	-1	1
Change in noncurrent operating liabilities	ncol_gr1a	Richardson et al. (2005)	1962-2001	-1	0
Change in net financial assets	nfna_gr1a	Richardson et al. (2005)	1962-2001	1	1
Earnings persistence	ni_ar1	Francis et al. (2004)	1975-2001	1	0
Net operating assets	noa_at	Hirshleifer et al. (2004)	1964-2002	-1	1

Investment

Liquidity of book assets	aliq_at	Ortiz-Molina and Phillips (2014)	1984-2006	-1	0
Asset Growth	at_gr1	Cooper Gulen and Schill (2008)	1968-2003	-1	1
Change in common equity	be_gr1a	Richardson et al. (2005)	1962-2001	-1	1
CAPEX growth (1 year)	capx_gr1	Xie (2001)	1971-1992	-1	0
CAPEX growth (2 years)	capx_gr2	Anderson and Garcia-Feijoo (2006)	1976-1998	-1	1
CAPEX growth (3 years)	capx_gr3	Anderson and Garcia-Feijoo (2006)	1976-1998	-1	1
Change in current operating assets	coa_gr1a	Richardson et al. (2005)	1962-2001	-1	1
Change in current operating liabilities	col_gr1a	Richardson et al. (2005)	1962-2001	-1	1
Hiring rate	emp_gr1	Belo Lin and Bazdresch (2014)	1965-2010	-1	1
Inventory growth	inv_gr1	Belo and Lin (2011)	1965-2009	-1	1
Inventory change	inv_gr1a	Thomas and Zhang (2002)	1970-1997	-1	1
Change in long-term net operating assets	lnoa_gr1a	Fairfield Whisenant and Yohn (2003)	1964-1993	-1	1
Mispricing factor: Management	mispricing_mgmt	Stambaugh and Yuan (2016)	1967-2013	1	1
Change in noncurrent operating assets	ncoa_gr1a	Richardson et al. (2005)	1962-2001	-1	1
Change in net noncurrent operating assets	nncoa_gr1a	Richardson et al. (2005)	1962-2001	-1	1
Change in net operating assets	noa_gr1a	Hirshleifer et al. (2004)	1964-2002	-1	1
Change PPE and Inventory	ppeinv_gr1a	Lyandres Sun and Zhang (2008)	1970-2005	-1	1
Long-term reversal	ret_60_12	De Bondt and Thaler (1985)	1926-1982	-1	1
Sales Growth (1 year)	sale_gr1	Lakonishok Shleifer and Vishny (1994)	1968-1989	-1	1
Sales Growth (3 years)	sale_gr3	Lakonishok Shleifer and Vishny (1994)	1968-1989	-1	1
Sales growth (1 quarter)	saleq_gr1		1967-2016	-1	0
Years 2-5 lagged returns, nonannual	seas_2.5na	Heston and Sadka (2008)	1965-2002	-1	1

Low Leverage

Firm age	age	Jiang Lee and Zhang (2005)	1965-2001	-1	1
Liquidity of market assets	aliq_mat	Ortiz-Molina and Phillips (2014)	1984-2006	-1	0
Book leverage	at_be	Fama and French (1992)	1963-1990	-1	0
The high-low bid-ask spread	bidaskhl_21d	Corwin and Schultz (2012)	1927-2006	1	1
Cash-to-assets	cash_at	Palazzo (2012)	1972-2009	1	0
Net debt-to-price	netdebt_me	Penman Richardson and Tuna (2007)	1962-2001	-1	1
Earnings volatility	ni_ivol	Francis et al. (2004)	1975-2001	1	0

R&D-to-sales	rd_sale	Chan Lakonishok and Sougiannis (2001)	1975-1995	1	0
R&D capital-to-book assets	rd5_at	Li (2011)	1952-2004	1	0
Asset tangibility	tangibility	Hahn and Lee (2009)	1973-2001	1	0
Altman Z-score	z_score	Dichev (1998)	1981-1995	1	1

Low Risk

Market Beta	beta_60m	Fama and MacBeth (1973)	1935-1968	-1	1
Dimson beta	beta_dimson_21d	Dimson (1979)	1955-1974	-1	0
Frazzini-Pedersen market beta	betabab_1260d	Frazzini and Pedersen (2014)	1926-2012	-1	1
Downside beta	betadown_252d	Ang Chen and Xing (2006)	1963-2001	-1	1
Earnings variability	earnings_variability	Francis et al. (2004)	1975-2001	-1	0
Idiosyncratic volatility from the CAPM (21 days)	ivol_capm_21d		1967-2016	-1	0
Idiosyncratic volatility from the CAPM (252 days)	ivol_capm_252d	Ali Hwang and Trombley (2003)	1976-1997	-1	1
Idiosyncratic volatility from the Fama-French 3-factor model	ivol_ff3_21d	Ang et al. (2006)	1963-2000	-1	1
Idiosyncratic volatility from the q-factor model	ivol_hxz4_21d		1967-2016	-1	0
Cash flow volatility	ocfq_saleq_std	Huang (2009)	1980-2004	-1	1
Maximum daily return	rmax1_21d	Bali Cakici and Whitelaw (2011)	1962-2005	-1	1
Highest 5 days of return	rmax5_21d	Bali, Brown, Murray and Tang (2017)	1993-2012	-1	1
Return volatility	rvol_21d	Ang et al. (2006)	1963-2000	-1	1
Years 6-10 lagged returns, nonannual	seas_6_10na	Heston and Sadka (2008)	1965-2002	-1	1
Share turnover	turnover_126d	Datar Naik and Radcliffe (1998)	1963-1991	-1	1
Number of zero trades with turnover as tiebreaker (1 month)	zero_trades_21d	Liu (2006)	1963-2003	1	0
Number of zero trades with turnover as tiebreaker (6 months)	zero_trades_126d	Liu (2006)	1963-2003	1	1
Number of zero trades with turnover as tiebreaker (12 months)	zero_trades_252d	Liu (2006)	1963-2003	1	1

Momentum

Current price to high price over last year	prc_highprc_252d	George and Hwang (2004)	1963-2001	1	1
Residual momentum t-6 to t-1	resff3_6_1	Blitz Huij and Martens (2011)	1930-2009	1	1
Residual momentum t-12 to t-1	resff3_12_1	Blitz Huij and Martens (2011)	1930-2009	1	1
Price momentum t-3 to t-1	ret_3_1	Jegadeesh and Titman (1993)	1965-1989	1	1
Price momentum t-6 to t-1	ret_6_1	Jegadeesh and Titman (1993)	1965-1989	1	1
Price momentum t-9 to t-1	ret_9_1	Jegadeesh and Titman (1993)	1965-1989	1	1
Price momentum t-12 to t-1	ret_12_1	Jegadeesh and Titman (1993)	1965-1989	1	1
Year 1-lagged return, nonannual	seas_1_1na	Heston and Sadka (2008)	1965-2002	1	1

Profit Growth

Change sales minus change Inventory	dsale_dinv	Abarbanell and Bushee (1998)	1974-1988	1	1
Change sales minus change receivables	dsale_drec	Abarbanell and Bushee (1998)	1974-1988	-1	0
Change sales minus change SG&A	dsale_dsga	Abarbanell and Bushee (1998)	1974-1988	1	0
Change in quarterly return on assets	niq_at_chg1		1972-2016	1	0

Change in quarterly return on equity	niq_be_chg1		1967-2016	1	0
Standardized earnings surprise	niq_su	Foster Olsen and Shevlin (1984)	1974-1981	1	1
Change in operating cash flow to assets	ocf_at_chg1	Bouchard, Krueger, Landier and Thesmar (2019)	1990-2015	1	1
Price momentum t-12 to t-7	ret_12_7	Novy-Marx (2012)	1925-2010	1	1
Labor force efficiency	sale_emp_gr1	Abarbanell and Bushee (1998)	1974-1988	1	0
Standardized Revenue surprise	saleq_su	Jegadeesh and Livnat (2006)	1987-2003	1	1
Year 1-lagged return, annual	seas_1_1an	Heston and Sadka (2008)	1965-2002	1	1
Tax expense surprise	tax_gr1a	Thomas and Zhang (2011)	1977-2006	1	1

Profitability

Coefficient of variation for dollar trading volume	dolvol_var_126d	Chordia Subrahmanyam and Anshuman (2001)	1966-1995	-1	1
Return on net operating assets	ebit_bev	Soliman (2008)	1984-2002	1	1
Profit margin	ebit_sale	Soliman (2008)	1984-2002	1	1
Pitroski F-score	f_score	Piotroski (2000)	1976-1996	1	1
Return on equity	ni_be	Haugen and Baker (1996)	1979-1993	1	1
Quarterly return on equity	niq_be	Hou Xue and Zhang (2015)	1972-2012	1	1
Ohlson O-score	o_score	Dichev (1998)	1981-1995	-1	1
Operating cash flow to assets	ocf_at	Bouchard, Krueger, Landier and Thesmar (2019)	1990-2015	1	1
Operating profits-to-book equity	ope_be	Fama and French (2015)	1963-2013	1	1
Operating profits-to-lagged book equity	ope_bell		1967-2016	1	0
Coefficient of variation for share turnover	turnover_var_126d	Chordia Subrahmanyam and Anshuman (2001)	1966-1995	-1	1

Quality

Capital turnover	at_turnover	Haugen and Baker (1996)	1979-1993	1	0
Cash-based operating profits-to-book assets	cop_at		1967-2016	1	0
Cash-based operating profits-to-lagged book assets	cop_atl1	Ball et al. (2016)	1963-2014	1	1
Change gross margin minus change sales	dgp_dsale	Abarbanell and Bushee (1998)	1974-1988	1	0
Gross profits-to-assets	gp_at	Novy-Marx (2013)	1963-2010	1	1
Gross profits-to-lagged assets	gp_atl1		1967-2016	1	0
Mispricing factor: Performance	mispricing_perf	Stambaugh and Yuan (2016)	1967-2013	1	1
Number of consecutive quarters with earnings increases	ni_inc8q	Barth Elliott and Finn (1999)	1982-1992	1	0
Quarterly return on assets	niq_at	Balakrishnan Bartov and Faurel (2010)	1976-2005	1	1
Operating profits-to-book assets	op_at		1963-2013	1	1
Operating profits-to-lagged book assets	op_atl1	Ball et al. (2016)	1963-2014	1	1
Operating leverage	opex_at	Novy-Marx (2011)	1963-2008	1	1
Quality minus Junk: Composite	qmj	Assness, Frazzini and Pedersen (2018)	1957-2016	1	1
Quality minus Junk: Growth	qmj_growth	Assness, Frazzini and Pedersen (2018)	1957-2016	1	1
Quality minus Junk: Profitability	qmj_prof	Assness, Frazzini and Pedersen (2018)	1957-2016	1	1

Quality minus Junk: Safety	qmj_safety	Assness, Frazzini and Pedersen (2018)	1957-2016	1	1
Assets turnover	sale_bev	Soliman (2008)	1984-2002	1	1
Seasonality					
Market correlation	corr_1260d	Assness, Frazzini, Gormsen, Pedersen (2020)	1925-2015	-1	1
Coskewness	coskew_21d	Harvey and Siddique (2000)	1963-1993	-1	1
Net debt issuance	dbnetis_at	Bradshaw Richardson and Sloan (2006)	1971-2000	-1	1
Kaplan-Zingales index	kz_index	Lamont Polk and Saa-Requejo (2001)	1968-1995	1	1
Change in long-term investments	lti_gr1a	Richardson et al. (2005)	1962-2001	-1	1
Taxable income-to-book income	pi_nix	Lev and Nissim (2004)	1973-2000	1	1
Years 2-5 lagged returns, annual	seas_2_5an	Heston and Sadka (2008)	1965-2002	1	1
Years 6-10 lagged returns, annual	seas_6_10an	Heston and Sadka (2008)	1965-2002	1	1
Years 11-15 lagged returns, annual	seas_11_15an	Heston and Sadka (2008)	1965-2002	1	1
Years 11-15 lagged returns, nonannual	seas_11_15na	Heston and Sadka (2008)	1965-2002	-1	0
Years 16-20 lagged returns, annual	seas_16_20an	Heston and Sadka (2008)	1965-2002	1	1
Change in short-term investments	sti_gr1a	Richardson et al. (2005)	1962-2001	1	0
Size					
Amihud Measure	ami_126d	Amihud (2002)	1964-1997	1	1
Dollar trading volume	dolvol_126d	Brennan Chordia and Subrahmanyam (1998)	1966-1995	-1	1
Market Equity	market_equity	Banz (1981)	1926-1975	-1	1
Price per share	prc	Miller and Scholes (1982)	1940-1978	-1	1
R&D-to-market	rd_me	Chan Lakonishok and Sougiannis (2001)	1975-1995	1	1
Skewness					
Idiosyncratic skewness from the CAPM	iskew_capm_21d		1967-2016	-1	0
Idiosyncratic skewness from the Fama-French 3-factor model	iskew_ff3_21d	Bali Engle and Murray (2016)	1925-2021	-1	1
Idiosyncratic skewness from the q-factor model	iskew_hxz4_21d		1967-2016	-1	0
Short-term reversal	ret_1.0	Jegadeesh (1990)	1929-1982	-1	1
Highest 5 days of return scaled by volatility	rmax5_rvol_21d	Assness, Frazzini, Gormsen, Pedersen (2020)	1925-2015	-1	1
Total skewness	rskew_21d	Bali Engle and Murray (2016)	1925-2021	-1	1
Value					
Assets-to-market	at_me	Fama and French (1992)	1963-1990	1	0
Book-to-market equity	be_me	Rosenberg Reid and Lanstein (1985)	1973-1984	1	1
Book-to-market enterprise value	bev_mev	Penman Richardson and Tuna (2007)	1962-2001	1	1
Net stock issues	chcsho_12m	Pontiff and Woodgate (2008)	1970-2003	-1	1
Debt-to-market	debt_me	Bhandari (1988)	1948-1979	1	1
Dividend yield	div12m_me	Litzenberger and Ramaswamy (1979)	1940-1980	1	1
Ebitda-to-market enterprise value	ebitda_mev	Loughran and Wellman (2011)	1963-2009	1	1

Equity duration	<u>eq_dur</u>	Dechow Sloan and Soliman (2004)	1962-1998	-1	1
Net equity issuance	<u>eqnetis_at</u>	Bradshaw Richardson and Sloan (2006)	1971-2000	-1	1
Equity net payout	<u>eqnpo_12m</u>	Daniel and Titman (2006)	1968-2003	1	1
Net payout yield	<u>eqnpo_me</u>	Boudoukh et al. (2007)	1984-2003	1	1
Payout yield	<u>eqpo_me</u>	Boudoukh et al. (2007)	1984-2003	1	1
Free cash flow-to-price	<u>fcf_me</u>	Lakonishok Shleifer and Vishny (1994)	1963-1990	1	1
Intrinsic value-to-market	<u>ival_me</u>	Frankel and Lee (1998)	1975-1993	1	0
Net total issuance	<u>netis_at</u>	Bradshaw Richardson and Sloan (2006)	1971-2000	-1	1
Earnings-to-price	<u>ni_me</u>	Basu (1983)	1963-1979	1	1
Operating cash flow-to-market	<u>ocf_me</u>	Desai Rajgopal and Venkatachalam (2004)	1973-1997	1	1
Sales-to-market	<u>sale_me</u>	Barbee Mukherji and Raines (1996)	1979-1991	1	1

Note: This table shows cluster names as underlined section headings and, for each cluster, a description of the factors included, the variable name used in the code, the original reference, the sample period used in the original reference, the sign of the factor (“1” means “long”, “-1” means “short”), and whether the original reference found the factor to be significant (“1” means “yes”, “0” means “no”). For example, the first value factor “at_me” goes long stocks with high values of assets-to-market and shorts those with low values (and would be done the reverse if the sign was “-1” instead of “1”).

Table J.2: Alpha Across Regions

Factor	US			Developed ex. US			Emerging			
	α_{OLS}	α_{EB}	$\Pr(\alpha_{EB} < 0)$	α_{OLS}	α_{EB}	$\Pr(\alpha_{EB} < 0)$	α_{OLS}	α_{EB}	$\Pr(\alpha_{EB} < 0)$	
1	aliq_mat*	-0.35	-0.31	1.00	-0.33	-0.27	1.00	-0.31	-0.28	1.00
2	dsale_drec*	-0.25	-0.22	1.00	-0.05	-0.11	0.86	-0.23	-0.18	0.95
3	bidaskhl_21d	-0.24	-0.25	1.00	-0.62	-0.42	1.00	-0.50	-0.38	1.00
4	ni_ivol*	-0.23	-0.18	0.99	-0.20	-0.13	0.90	0.03	-0.07	0.75
5	at_be*	-0.16	-0.10	0.92	0.06	0.04	0.34	0.12	0.04	0.37
6	age	-0.15	-0.18	1.00	-0.64	-0.44	1.00	-0.59	-0.41	1.00
7	kz_index	-0.13	-0.12	0.93	-0.08	-0.08	0.79	-0.32	-0.15	0.93
8	turnover_var_126d	-0.13	-0.12	0.96	-0.02	-0.03	0.60	0.20	0.01	0.46
9	prc	-0.11	-0.02	0.60	0.02	0.06	0.26	0.07	0.07	0.23
10	sti_gr1a*	-0.09	-0.02	0.60	-0.01	0.03	0.37	0.17	0.08	0.21
11	dolvol_var_126d	-0.07	-0.07	0.85	-0.07	-0.03	0.63	0.20	0.02	0.43
12	dsale_dsga*	-0.07	-0.02	0.58	0.20	0.13	0.11	0.33	0.15	0.09
13	ni_ar1*	-0.02	-0.07	0.81	0.04	0.03	0.39	-0.29	-0.06	0.71
14	sale_emp_gr1*	-0.01	-0.03	0.64	-0.24	-0.10	0.83	0.11	-0.00	0.51
15	netdebt_me	-0.01	0.03	0.32	0.11	0.11	0.13	0.20	0.12	0.13
16	z_score	-0.00	0.03	0.36	-0.02	0.03	0.36	0.20	0.10	0.17
17	iskew_hxz4_21d*	0.01	-0.08	0.80	-0.50	-0.18	0.94	-0.37	-0.16	0.90
18	rd_sale*	0.01	0.06	0.22	0.22	0.18	0.03	0.25	0.15	0.08
19	market_equity	0.02	0.13	0.04	0.13	0.21	0.02	0.55	0.37	0.00
20	cash_at*	0.04	0.07	0.18	0.07	0.10	0.15	0.24	0.14	0.08
21	ami_126d	0.05	0.14	0.03	0.14	0.21	0.02	0.38	0.28	0.00
22	ncol_gr1a*	0.05	-0.01	0.57	-0.07	0.02	0.42	-0.08	0.02	0.41
23	debt_me	0.09	0.06	0.21	0.03	0.05	0.31	-0.06	-0.00	0.51
24	ni_inc8q*	0.10	0.13	0.07	0.38	0.25	0.02	0.26	0.18	0.07
25	tax_gr1a	0.10	0.09	0.13	0.06	0.10	0.17	0.30	0.16	0.06
26	saleq_gr1*	0.11	-0.03	0.66	0.03	0.04	0.37	-0.70	-0.16	0.91
27	ret_60_12	0.12	0.03	0.31	0.11	0.19	0.03	0.37	0.30	0.00
28	rd5_at*	0.12	0.17	0.01	0.35	0.29	0.00	0.63	0.31	0.01
29	coskew_21d	0.12	0.12	0.05	0.28	0.19	0.02	-0.04	0.08	0.20
30	saleq_su	0.12	0.16	0.04	0.47	0.30	0.01	0.54	0.29	0.01
31	col_gr1a	0.13	0.01	0.46	0.00	0.07	0.23	-0.11	0.05	0.31
32	iskew_ff3_21d	0.13	0.08	0.18	-0.20	0.01	0.46	0.15	0.11	0.17
33	tangibility*	0.13	0.17	0.01	0.26	0.25	0.00	0.41	0.28	0.00
34	lti_gr1a	0.15	0.09	0.13	-0.01	0.04	0.35	-0.17	-0.01	0.52
35	pi_nix	0.16	0.14	0.05	0.06	0.11	0.14	0.09	0.12	0.13
36	bev_mev	0.17	0.16	0.02	0.28	0.25	0.00	0.27	0.22	0.02
37	gp_atl1*	0.18	0.20	0.01	0.22	0.24	0.01	0.53	0.31	0.00
38	at_me*	0.18	0.16	0.01	0.20	0.20	0.02	0.23	0.19	0.04
39	seas_16_20na	0.19	0.04	0.32	-0.13	0.02	0.44	-0.32	-0.02	0.57
40	zero_trades_21d*	0.19	0.16	0.01	0.21	0.18	0.04	0.40	0.26	0.01
41	ebit_sale	0.20	0.18	0.01	0.21	0.18	0.03	0.25	0.15	0.07
42	ret_3_1	0.21	0.11	0.06	0.26	0.15	0.06	0.08	0.08	0.21
43	be_me	0.23	0.22	0.00	0.32	0.29	0.00	0.32	0.27	0.01
44	op_atl1	0.23	0.23	0.00	0.25	0.25	0.01	0.43	0.27	0.00
45	at_turnover*	0.23	0.23	0.00	0.23	0.24	0.01	0.37	0.25	0.01
46	opex_at	0.24	0.22	0.00	0.22	0.22	0.01	0.23	0.18	0.04
47	ope_bell*	0.24	0.23	0.00	0.30	0.27	0.00	0.43	0.26	0.00
48	earnings_variability*	0.24	0.16	0.02	0.09	0.07	0.24	-0.00	0.06	0.28
49	seas_1_1na	0.26	0.17	0.01	0.23	0.14	0.06	0.28	0.21	0.02
50	dolvol_126d	0.26	0.34	0.00	0.27	0.35	0.00	0.53	0.43	0.00
51	be_gr1a	0.26	0.13	0.05	0.09	0.15	0.06	-0.10	0.10	0.16
52	div12m_me	0.27	0.27	0.00	0.51	0.46	0.00	0.71	0.51	0.00
53	sale_me	0.28	0.26	0.00	0.35	0.33	0.00	0.42	0.34	0.00

54	ocfq_saleq_std	0.28	0.25	0.00	0.38	0.30	0.00	0.69	0.39	0.00
55	niq_at	0.29	0.36	0.00	0.75	0.57	0.00	0.96	0.58	0.00
56	sale_gr3	0.29	0.16	0.02	0.09	0.19	0.03	0.15	0.23	0.02
57	sale_gr1	0.30	0.16	0.02	0.17	0.20	0.02	-0.19	0.08	0.22
58	ni_be	0.30	0.28	0.00	0.42	0.33	0.00	0.31	0.23	0.01
59	ivol_capm_252d	0.30	0.26	0.00	0.41	0.30	0.00	0.21	0.23	0.01
60	seas_2_5na	0.31	0.22	0.00	0.36	0.41	0.00	0.67	0.54	0.00
61	ret_6_1	0.31	0.22	0.00	0.33	0.23	0.01	0.38	0.29	0.00
62	seas_11_15na*	0.31	0.22	0.00	-0.07	0.08	0.23	-0.26	0.05	0.34
63	o_score	0.31	0.29	0.00	0.38	0.33	0.00	0.45	0.31	0.00
64	beta_dimson_21d*	0.31	0.26	0.00	0.33	0.24	0.00	0.06	0.14	0.08
65	aliq_at*	0.31	0.17	0.01	0.04	0.14	0.08	-0.07	0.12	0.13
66	dgp_dsale*	0.31	0.23	0.00	-0.11	0.02	0.44	-0.03	0.02	0.43
67	rd_me	0.32	0.39	0.00	0.36	0.39	0.00	0.39	0.39	0.00
68	corr_1260d	0.32	0.29	0.00	0.24	0.27	0.00	0.26	0.26	0.00
69	qmj_safety	0.33	0.34	0.00	0.32	0.36	0.00	0.81	0.51	0.00
70	emp_gr1	0.34	0.22	0.00	0.16	0.28	0.00	0.50	0.41	0.00
71	eq_dur	0.34	0.33	0.00	0.38	0.37	0.00	0.52	0.41	0.00
72	betadown_252d	0.35	0.31	0.00	0.48	0.36	0.00	0.25	0.27	0.00
73	prc_highprc_252d	0.35	0.26	0.00	0.38	0.27	0.00	0.40	0.32	0.00
74	turnover_126d	0.36	0.32	0.00	0.50	0.39	0.00	0.50	0.39	0.00
75	ret_1_0	0.36	0.29	0.00	0.20	0.28	0.00	0.07	0.20	0.03
76	gp_at	0.36	0.37	0.00	0.31	0.36	0.00	0.78	0.49	0.00
77	beta_60m	0.37	0.32	0.00	0.38	0.32	0.00	0.43	0.35	0.00
78	taccruals_at	0.37	0.20	0.01	0.06	0.15	0.07	-0.00	0.12	0.13
79	zero_trades_126d	0.38	0.34	0.00	0.52	0.41	0.00	0.51	0.41	0.00
80	ival_me*	0.38	0.35	0.00	0.38	0.37	0.00	0.40	0.36	0.00
81	seas_2_5an	0.39	0.38	0.00	0.55	0.45	0.00	0.46	0.41	0.00
82	ope_be	0.39	0.36	0.00	0.45	0.39	0.00	0.49	0.36	0.00
83	sale_bev	0.40	0.39	0.00	0.32	0.35	0.00	0.59	0.41	0.00
84	niq_be	0.41	0.43	0.00	0.83	0.61	0.00	0.92	0.59	0.00
85	niq_at_chg1*	0.41	0.38	0.00	0.18	0.32	0.00	0.90	0.48	0.00
86	ret_9_1	0.41	0.32	0.00	0.42	0.32	0.00	0.57	0.44	0.00
87	taccruals_ni	0.42	0.22	0.00	0.04	0.12	0.11	-0.18	0.04	0.34
88	dbnetis_at	0.42	0.37	0.00	0.14	0.28	0.00	0.59	0.41	0.00
89	ebit_bev	0.42	0.39	0.00	0.33	0.35	0.00	0.62	0.41	0.00
90	iskew_capm_21d*	0.42	0.33	0.00	-0.04	0.16	0.05	0.21	0.24	0.01
91	eqpo_me	0.42	0.38	0.00	0.37	0.37	0.00	0.43	0.37	0.00
92	seas_16_20an	0.42	0.38	0.00	0.51	0.38	0.00	0.22	0.33	0.00
93	op_at	0.42	0.40	0.00	0.46	0.42	0.00	0.52	0.40	0.00
94	betabab_1260d	0.42	0.39	0.00	0.61	0.50	0.00	0.70	0.53	0.00
95	seas_1_1an	0.42	0.36	0.00	0.47	0.36	0.00	0.11	0.23	0.01
96	ivol_capm_21d*	0.42	0.38	0.00	0.53	0.43	0.00	0.45	0.40	0.00
97	ni_me	0.43	0.41	0.00	0.41	0.43	0.00	0.76	0.55	0.00
98	seas_11_15an	0.43	0.38	0.00	0.20	0.31	0.00	0.39	0.34	0.00
99	at_gr1	0.43	0.29	0.00	0.08	0.22	0.01	0.24	0.31	0.00
100	zero_trades_252d	0.44	0.40	0.00	0.57	0.45	0.00	0.54	0.45	0.00
101	ivol_hxz4_21d*	0.44	0.39	0.00	0.38	0.37	0.00	0.88	0.52	0.00
102	ebitda_mev	0.44	0.43	0.00	0.46	0.46	0.00	0.70	0.53	0.00
103	eqnpo_12m	0.45	0.44	0.00	0.64	0.58	0.00	0.72	0.58	0.00
104	capx_gr3	0.45	0.31	0.00	0.41	0.39	0.00	-0.02	0.25	0.01
105	niq_su	0.47	0.41	0.00	0.24	0.35	0.00	0.78	0.46	0.00
106	rvol_21d	0.47	0.42	0.00	0.52	0.42	0.00	0.33	0.35	0.00
107	ivol_ff3_21d	0.48	0.43	0.00	0.45	0.41	0.00	0.72	0.52	0.00
108	ocf_me	0.49	0.46	0.00	0.31	0.40	0.00	0.92	0.63	0.00
109	ret_12_7	0.49	0.44	0.00	0.54	0.45	0.00	0.41	0.39	0.00

110	capex_abn	0.49	0.35	0.00	0.27	0.32	0.00	0.06	0.27	0.01
111	niq_be_chg1*	0.50	0.44	0.00	0.20	0.35	0.00	0.87	0.50	0.00
112	chcsho_12m	0.50	0.46	0.00	0.42	0.41	0.00	0.28	0.33	0.00
113	resff3_6_1	0.51	0.39	0.00	0.56	0.43	0.00	0.47	0.41	0.00
114	eqnpo_me	0.51	0.47	0.00	0.35	0.40	0.00	0.62	0.49	0.00
115	coa_gr1a	0.51	0.37	0.00	0.17	0.30	0.00	0.26	0.35	0.00
116	qmj_growth	0.52	0.44	0.00	0.24	0.29	0.00	0.25	0.26	0.01
117	eqnetis_at	0.52	0.49	0.00	0.54	0.52	0.00	0.59	0.51	0.00
118	ret_12_1	0.52	0.42	0.00	0.47	0.37	0.00	0.47	0.41	0.00
119	rskew_21d	0.52	0.42	0.00	0.08	0.26	0.01	0.19	0.27	0.00
120	netis_at	0.53	0.50	0.00	0.53	0.52	0.00	0.68	0.55	0.00
121	seas_6_10na	0.53	0.47	0.00	0.65	0.50	0.00	0.36	0.40	0.00
122	mispricing_perf	0.55	0.52	0.00	0.57	0.51	0.00	0.53	0.46	0.00
123	f_score	0.56	0.51	0.00	0.51	0.47	0.00	0.57	0.45	0.00
124	rmax1_21d	0.56	0.50	0.00	0.59	0.48	0.00	0.38	0.40	0.00
125	ocf_at_chg1	0.56	0.49	0.00	0.48	0.45	0.00	0.46	0.43	0.00
126	seas_6_10an	0.57	0.54	0.00	0.99	0.65	0.00	0.33	0.48	0.00
127	fcf_me	0.58	0.56	0.00	0.50	0.55	0.00	1.05	0.73	0.00
128	rmax5_21d	0.58	0.51	0.00	0.63	0.51	0.00	0.31	0.38	0.00
129	capx_gr2	0.60	0.44	0.00	0.35	0.41	0.00	0.10	0.34	0.00
130	qmj_prof	0.60	0.55	0.00	0.48	0.48	0.00	0.59	0.48	0.00
131	qmj	0.60	0.56	0.00	0.55	0.52	0.00	0.62	0.50	0.00
132	capx_gr1*	0.64	0.47	0.00	0.32	0.41	0.00	0.15	0.37	0.00
133	debt_gr3	0.67	0.50	0.00	0.36	0.43	0.00	0.22	0.39	0.00
134	rmax5_rvol_21d	0.67	0.56	0.00	0.34	0.43	0.00	0.04	0.29	0.00
135	lnoa_gr1a	0.69	0.52	0.00	0.29	0.39	0.00	0.07	0.33	0.00
136	fml_gr1a	0.71	0.54	0.00	0.25	0.40	0.00	0.45	0.48	0.00
137	ppeinv_gr1a	0.71	0.52	0.00	0.24	0.36	0.00	0.07	0.32	0.00
138	oaccruals_at	0.71	0.54	0.00	0.38	0.48	0.00	0.65	0.57	0.00
139	inv_gr1a	0.72	0.55	0.00	0.25	0.39	0.00	0.29	0.44	0.00
140	nfna_gr1a	0.73	0.56	0.00	0.20	0.40	0.00	0.56	0.53	0.00
141	cop_at11	0.75	0.68	0.00	0.46	0.52	0.00	0.77	0.61	0.00
142	dsale_dinv	0.75	0.57	0.00	0.21	0.30	0.00	0.03	0.25	0.01
143	oaccruals_ni	0.76	0.55	0.00	0.15	0.35	0.00	0.57	0.51	0.00
144	ocf_at	0.77	0.71	0.00	0.72	0.65	0.00	0.68	0.58	0.00
145	mispricing_mgmt	0.78	0.63	0.00	0.58	0.64	0.00	0.66	0.68	0.00
146	ncoa_gr1a	0.80	0.61	0.00	0.33	0.45	0.00	0.29	0.47	0.00
147	inv_gr1	0.80	0.62	0.00	0.47	0.52	0.00	0.06	0.38	0.00
148	resff3_12_1	0.81	0.69	0.00	0.98	0.80	0.00	0.72	0.68	0.00
149	nncoa_gr1a	0.82	0.63	0.00	0.28	0.43	0.00	0.30	0.47	0.00
150	cowc_gr1a	0.85	0.64	0.00	0.40	0.51	0.00	0.57	0.57	0.00
151	noa_at	0.87	0.66	0.00	0.36	0.47	0.00	0.22	0.44	0.00
152	noa_gr1a	0.91	0.75	0.00	0.54	0.64	0.00	0.64	0.70	0.00
153	cop_at*	1.02	0.91	0.00	0.62	0.69	0.00	0.85	0.73	0.00

Note: The table shows monthly alpha in percentages across three different regions. α_{OLS} is the intercept from an OLS regression of the factor return on the regional market return. α_{EB} is the factor-region specific posterior mean found via the empirical Bayes procedure applied jointly to all the factor-region specific factors. $\Pr(\alpha_{EB} < 0)$ is the probability that the alpha is negative based on the posterior distribution from the EB procedure. We count a factor as replicated if this probability is below 2.5%. The residual volatility of all strategies have been scaled to 10% annualized. A “*” indicates that the original paper did not propose the factors as a significant predictor of realized returns.

Table J.3: Country Information

	Country	MSCI	Start	Stocks	Mega Stocks	Total Market Cap	Median MC
1	USA	Developed	1926-01-31	4,356	449	4.28e+07	796
2	CHN	Emerging	1991-02-28	4,007	102	1.07e+07	816
3	JPN	Developed	1986-01-31	3,867	88	7.05e+06	176
4	HKG	Developed	1986-01-31	2,352	61	5.33e+06	92
5	GBR	Developed	1985-12-31	1,650	36	3.14e+06	158
6	FRA	Developed	1986-01-31	693	36	2.94e+06	121
7	SAU	Emerging	2000-02-29	199	8	2.61e+06	460
8	DEU	Developed	1986-01-31	659	36	2.58e+06	155
9	IND	Emerging	1988-09-30	3,527	26	2.57e+06	10
10	KOR	Emerging	1986-02-28	2,258	21	2.09e+06	131
11	CAN	Developed	1982-03-31	718	35	2.08e+06	198
12	TWN	Emerging	1988-02-29	1,951	12	1.77e+06	121
13	AUS	Developed	1985-11-30	1,702	19	1.68e+06	51
14	CHE	Developed	1986-01-31	233	22	1.66e+06	863
15	NLD	Developed	1986-01-31	119	17	1.11e+06	939
16	SWE	Developed	1986-01-31	688	16	1.00e+06	82
17	ITA	Developed	1986-01-31	359	11	7.68e+05	121
18	ESP	Developed	1986-01-31	182	12	7.02e+05	239
19	RUS	Emerging	1995-08-31	186	11	6.48e+05	184
20	DNK	Developed	1986-01-31	161	8	5.90e+05	184
21	BRA	Emerging	1988-05-31	220	6	5.90e+05	575
22	SGP	Developed	1986-01-31	562	7	5.31e+05	59
23	THA	Emerging	1986-07-31	762	5	5.30e+05	75
24	IDN	Emerging	1989-01-31	656	6	4.93e+05	82
25	MYS	Emerging	1986-01-31	923	2	4.44e+05	55
26	ZAF	Emerging	1986-01-31	261	3	4.07e+05	137
27	BEL	Developed	1986-01-31	131	4	3.94e+05	431
28	NOR	Developed	1986-01-31	272	3	3.44e+05	230
29	MEX	Emerging	1986-02-28	117	3	3.39e+05	992
30	FIN	Developed	1986-01-31	152	6	3.15e+05	218
31	PHL	Emerging	1986-01-31	248	2	2.66e+05	164
32	TUR	Emerging	1990-03-31	389	1	2.38e+05	119
33	ARE	Emerging	2001-06-30	104	4	2.29e+05	284
34	ISR	Developed	1994-12-31	415	1	2.23e+05	110
35	VNM	Frontier	2006-08-31	713	0	1.80e+05	18
36	POL	Emerging	1993-07-31	706	1	1.78e+05	14
37	CHL	Emerging	1989-01-31	174	0	1.71e+05	221
38	QAT	Emerging	2001-12-31	47	2	1.65e+05	971
39	IRL	Developed	1986-01-31	33	4	1.65e+05	652
40	NZL	Developed	1986-01-31	122	0	1.38e+05	196
41	AUT	Developed	1986-01-31	68	0	1.26e+05	327
42	KWT	Frontier	2001-01-31	163	2	9.84e+04	80
43	PER	Emerging	1990-01-31	103	0	8.95e+04	81
44	COL	Emerging	1989-01-31	44	1	8.81e+04	404
45	PRT	Developed	1986-08-31	44	2	8.55e+04	97
46	MAR	Frontier	1995-09-30	71	0	6.51e+04	183
47	GRC	Emerging	1988-09-30	149	0	5.20e+04	33
48	PAK	Emerging	1992-09-30	411	0	4.97e+04	19
49	ARG	Emerging	1988-09-30	68	1	4.95e+04	113
50	NGA	Frontier	1993-11-30	150	0	4.68e+04	11
51	BGD	Frontier	2002-05-31	324	0	4.40e+04	30
52	EGY	Emerging	1996-12-31	204	0	3.87e+04	46

53	HUN	Emerging	1993-06-30	33	0	2.70e+04	54
54	CZE	Emerging	1995-01-31	14	0	2.67e+04	282
55	ROU	Frontier	1997-11-30	76	0	2.55e+04	29
56	BHR	Frontier	2001-03-31	26	0	2.14e+04	243
57	KEN	Frontier	1993-11-30	50	0	2.09e+04	30
58	HRV	Frontier	1997-11-30	73	0	2.07e+04	35
59	BGR	Standalone	1995-12-31	113	0	1.68e+04	24
60	JOR	Frontier	1993-08-31	154	0	1.63e+04	18
61	OMN	Frontier	1998-03-31	107	0	1.62e+04	35
62	LKA	Frontier	1987-06-30	265	0	1.55e+04	15
63	TTO	Standalone	1997-08-31	19	0	1.29e+04	292
64	KAZ	Frontier	2009-06-30	12	0	1.21e+04	394
65	ISL	Standalone	1995-12-31	22	0	1.20e+04	329
66	JAM	Standalone	1993-12-31	66	0	1.13e+04	29
67	SVN	Frontier	1995-03-31	22	0	8.50e+03	122
68	TUN	Frontier	1995-09-30	74	0	8.46e+03	38
69	CIV	Frontier	2002-04-30	39	0	7.61e+03	85
70	MUS	Frontier	1995-08-31	62	0	7.14e+03	48
71	LUX	Not Rated	1986-01-31	8	0	7.14e+03	393
72	LTU	Frontier	1995-10-31	28	0	5.51e+03	61
73	MLT	Standalone	1995-08-31	22	0	5.05e+03	124
74	LBN	Standalone	1997-11-30	8	0	4.12e+03	335
75	EST	Frontier	1996-01-31	19	0	3.48e+03	69
76	TZA	Not Rated	2000-07-31	16	0	3.34e+03	16
77	SRB	Frontier	2009-09-30	29	0	3.17e+03	10
78	BWA	Standalone	1995-09-30	22	0	3.13e+03	91
79	SVK	Not Rated	1986-01-31	10	0	3.12e+03	99
80	CYP	Not Rated	1994-01-31	37	0	3.12e+03	21
81	PSE	Standalone	2008-07-31	27	0	2.97e+03	59
82	GHA	Not Rated	1997-11-30	16	0	2.89e+03	67
83	BMU	Not Rated	2007-08-31	8	0	2.51e+03	118
84	NAM	Not Rated	1996-06-30	8	0	2.18e+03	329
85	MWI	Not Rated	2008-08-31	12	0	2.15e+03	118
86	ECU	Not Rated	1999-04-30	2	0	1.85e+03	927
87	LVA	Not Rated	1997-10-31	20	0	1.17e+03	13
88	UGA	Not Rated	2011-10-31	9	0	1.15e+03	97
89	ZMB	Not Rated	1996-03-31	10	0	5.27e+02	26
90	UKR	Standalone	2008-02-29	4	0	3.28e+02	52
91	GGY	Not Rated	2015-04-30	2	0	2.25e+02	113
92	IRN	Not Rated	2002-05-31	0	0	0.00e+00	0
93	URY	Not Rated	1996-06-30	0	0	0.00e+00	0
All				4.021700e+04	1,092	1.01e+08	

Note: The table shows summary statistics by the country where a security is listed. We include common stocks that are the primary security of the underlying firm, traded on a standard exchange, with non-missing return and market equity data. *Country* is the ISO code of the underlying exchange country. For further information, see https://en.wikipedia.org/wiki/List_of_ISO_3166_country_codes. *MSCI* shows the MSCI classification of each country as of January 7th 2021. For the most recent classification, see <https://www.msci.com/market-classification>. *Start* is the first date with a valid observation. In the next 4 columns, the data is shown as of December 31st 2020. *Stocks* is the number of stocks available. *Mega stocks* is the number of stocks with a market cap above the 80th percentile of NYSE stocks. *Total Market Cap* is the aggregate market cap in million USD. *Median MC* is the median market cap in million USD.

Data Documentation

K Global Factor Data Documentation

We end the Internet Appendix with a documentation of our global factor data and how to use it to replicate our results and for future research. We will continue to update this data and its documentation as seen on our websites. The online document also contains instructions on how to run the code, bug fixes, and so on.

K.1 Identifier Variables

Here we define important identifying variables for our empirical analysis. We assign stocks to countries by `excntry`. We assign stocks to size groups via `size_grp`. We only include stocks with 1 on all the the `obs_main`, `exch_main`, `primary_sec` and `common` indicators.

Table K.1: Identifier Variables

Name	Description
<code>excntry</code>	The country of the exchange where the security is traded. Usually expressed as an ISO currency code with the exception of <i>mul</i> which indicates a multi country exchange ⁵⁴
<code>size_grp</code>	This groups each firm into one of five categories: Mega, Large, Small, Micro and Nano cap. The groups are non-overlapping and the breakpoints are based on the market equity of NYSE stocks. In particular, Mega caps are all stocks with market equity larger than the 80th percentile of NYSE stocks, Large caps are all remaining stocks larger than the 50th percentile, Small caps are larger than the 20th percentile, Micro caps are larger than the 1st percentile and Nano caps are the remaining stocks.
<code>obs_main</code>	If there are more than one firm observations for one date, this identifies if the observation is considered as the 'main' observation. If available, CRSP observations are considered as the 'main' observation.
<code>exch_main</code>	Indicator for main exchanges. If CRSP is the source, main exchanges are those with <i>crsp_exchcd</i> 1, 2 and 3. If Compustat is the source, main exchanges are all <i>comp_exchg</i> except 0, 1, 2, 3, 4, 13, 15, 16, 17, 18, 19, 20, 21, 127, 150, 157, 229, 263, 269, 281, 283, 290, 320, 326, 341, 342, 347, 348, 349, 352.
<code>primary_sec</code>	Primary security as identified by Compustat. A 'gvkey' can have up to three different primary securities ('iid') at a given time (US, CA, and international). All observations from CRSP has <code>primary_sec=1</code> .
<code>common</code>	Indicator for common stocks. If CRSP is the source, common is one if the SHRCD variable is 10, 11 or 12. If Compustat is the source, common is one if TPCI is '0'

K.2 Helper Functions

This section describes functions that we use to create variables. Many of the functions are used for variables with quarterly, monthly and daily frequencies, and these are specified by “_zQ”, “_zM” and “_zD” respectively, where “z” is the number of quarters, months or days that the function is referencing. For example, `COVAR_12M(X, Y)` is the covariance of variables X and Y over the past 12 months.

Table K.2: Helper Functions

Function	Name	Description
Mean	\bar{X}_z	$\frac{1}{z} \sum_{n=0}^{z-1} X_{t-n}$
Variance	<code>VARC_z(X)</code>	$\frac{1}{z-1} \sum_{n=0}^{z-1} (X_{t-n} - \bar{X}_{tz})^2$
Covariance	<code>COVAR_z(X, Y)</code>	$\frac{1}{z-1} \sum_{n=0}^{z-1} (X_{t-n} - \bar{X}_{tz})(Y_{t-n} - \bar{Y}_{tz})$
Standard Deviation	$\sigma_z(X)$	$\sqrt{VARC_z(X)}$
Skewness	<code>SKEW_z(X)</code>	$\frac{1}{z \times \sigma_z(X)^3} \sum_{n=0}^{z-1} (X_{t-n} - \bar{X}_{tz})^3$

⁵⁴Typically over the counter exchanges.

Function	Name	Description
Standardized Unexpected Realization	SUR _z (X)	$\frac{X_t - (X_{t-3} + \overline{(X_{t-3} - X_{t-15})_z} / 4)}{\sigma_z(X_{t-3} - X_{t-15})}$
Change to Expectations	CHG_TO_EXP(X)	$\frac{X_t}{(X_{t-12} + X_{t-24}) / 2}$
Maximum	MAXn _z (X)	The maximum n values of given input.
Quality Minus Junk Helpers		
Earnings Volatility	<i>_EVOL</i>	<i>ROEQ_BE_STD</i> × 2. If this is unavailable, we use <i>ROE_BE_STD</i> .
Rank of Variable	<i>_rVar</i>	Cross-sectional rank of Var within a country ⁵⁵
Z transformation	<i>ZV(rVar)</i>	$\frac{-rVAR - \overline{-rVAR}_z}{.t(-rVAR)}$

K.3 Accounting Characteristics

Datasets

- COMP.FUNDA
- COMP.FUNDQ
- COMP.G_FUNDA
- COMP.G_FUNDQ

General Information

- We create characteristics for annual and quarterly accounting data separately. We then take the most recent characteristics value from each dataset to create the final dataset.
- We assume that accounting variables are publically available 4 months after the end of the accounting period.
- In describing accounting variables, we use the Compustat item names from the annual dataset. The equivalent item name in the quarterly dataset can be found by adding a ‘q’ or ‘y’ to the end of the annual item name. Specifically, ‘q’ indicates a value calculated over one quarter while ‘y’ refers to the cummulative value over the quarters with data available within a fiscal year.

Annualized Accounting Variables from Quarterly Data

- The value of a balance sheet item such as asset or book equity has the same meaning in the annual and the quarterly data. It is the value by the end of a fiscal period.
- The value of an income or cash flow statement item is different. In the annual data, it is calculated over one year. However, in the quarterly data, it is calculated over one quarter. To make quarterly income and cash flows items comparable to the corresponding annual item, we take the sum of the item over the last four quarters.

⁵⁵*OACCRUALS_AT*, *BETABAB_1260d*, *DEBT_AT* and *_EVOL* are sorted in descending order. All other variables are sorted in ascending order.

Accounting Variables

The abbreviation is used to refer to the accounting variable. A suffix of '*' indicates that we have altered the original Compustat item to increase the coverage or to create a variable that is a part of creating a characteristic in the final dataset. The characteristic name will reflect the accounting name except the '*' suffix. As an example, 'gp_at' is gross profit scaled by assets. In general, we will refer to Compustat variables using capital letters.

Table K.3: Accounting Variables

Name	Abbreviation	Construction
Income Statement		
Sales	sale*	We prefer SALE . If this is unavailable, we use REVT
Cost of Goods Sold	cogs	Compustat item COGS
Gross Profit	gp*	We prefer to use GP . If this is unavailable we use sale*- COGS
Selling, General and Administrative Expenses	xsga	Compustat item XSGA
Advertising Expenses	xad	Compustat item XAD . Note that this is not available in Compustat Global
Research and Development Expenses	xrd	Compustat item XRD . Note that this is not available in Compustat Global
Staff Expenses	xlr	Compustat item XLR
Special Items	spi	Compustat item SPI
Operating Expenses	opex*	We prefer to use XOPR . If this is unavailable, we use COGS+XSGA
Operating Income Before Depreciation	ebitda*	We prefer to use EBITDA . If this is unavailable, we use OIBDP . If this is unavailable, we use SALE*- OPEX* . If this is unavailable, we use GP*- XSGA
Depreciation and Amortization	dp	Compustat Item DP
Operating Income After Depreciation	ebit*	We prefer to use EBIT . If this is unavailable, we use OIADP . If this is unavailable, we use EBITDA*-DP
Interest Expenses	int	Compustat item XINT
Operating Profit ala Ball et al (2015)	op*	We use EBITDA* + XRD . If XRD is unavailable, we set it to zero
Operating Profit to Equity	ope*	We use EBITDA*-XINT . Note that we target the same variable as the numerator of the profitability characteristic used to create the Robust-minus weak factor in the fama-French 5 factor model (Fama and French, 2015)
Earnings before Tax and Extraordinary Items	pi*	We prefer to use PI . If this is unavailable we use EBIT*-XINT+SPI+NOPI where we set SPI and NOPI to zero if missing
Income Tax	tax	Compustat item TXT
Extraordinary Items and Discontinued Operations	xido*	We prefer to use XIDO . If this is unavailable, we use XI+DO where we set DO to zero if missing. The reason why we set missing DO to zero is because it is not available in COMP.G.FUNDQ
Net Income	ni*	We prefer to use IB . If this is unavailable, we use NI-XIDO* . If this is unavailable, we prefer PI*-TXT-MII . If MII is unavailable, it is set to zero
Net Income Including Extraordinary Items	nix*	We prefer NI . If this is not available, we prefer NI*+XIDO* . If XIDO* is unavailable, we set it to zero. If that is unavailable, we prefer NI*+XI+DO
Firm Income	fi*	We use NIX*+XINT
Dividends for Common Shareholds	dvc	Compustat Item DVC
Total Dividends	div*	We prefer DVT . If this is not available, we use DV
Income Before Extraordinary Items	ni_qtr*	We use IBQ
Net Sales	sale_qtr*	We use SALEQ
Cash Flow Statement		
Capital Expenditures	capx	Compustat item CAPX
Capital Expenditures to Sales	capex_sale*	We use CAPX / SALE*
Free Cash Flow	fcf*	We use OCF*-CAPX . Note that the free cash flow is computed before financing activities and sale of assets is taken into account
Equity Buyback	eqbb*	We use PRSTKC+PURTSHR Equity Buyback is mainly PRSTKC in NA and PURTSHR in GLOBAL. Either of PRSTKC or PURTSHR are allowed to be missing
Equity Issuance	eqis*	Compustat item SSTK

Name	Abbreviation	Construction
Equity Net Issuance	eqnetis*	We use EQIS*-EQBB*. Either EQIS* or EQBB* are allowed to be missing
Net Equity Payout	eqpo*	We use DIV*+EQBB*
Equity Net Payout	eqnpo*	We use DIV*-EQNETIS*
Net Long-Term Debt Issuance	dltnetis*	We prefer to use DLTIS-DLTR where we only require that one of the items are non-missing. If this is unavailable, we use LTDCH . If this is unavailable we use the yearly change in long-term book debt DLTT
Net Short-Term Debt Issuance	dstnetis*	We prefer DLCCH . If this is unavailable, we use the yearly change in short-term book debt DLC
Net Debt Issuance	dbnetis*	We use DLTNETIS*+DSTNETIS* and only require one of the items to be non-missing
Net Issuance	netis*	We use EQNETIS*+DBNETIS*. Either EQNETIS* or DBNETIS* are allowed to be missing
Financial Cash Flow	fincf*	We prefer FINCF . If this is unavailable, we use NETIS*- DV + FIAO + TXBCOF . If FIAO or TXBCOF is missing, it is set to zero
Balance Sheet - Assets		
Total Assets	at*	We prefer to use AT . If this is unavailable, then we use SEQ* + DLTT + LCT + LO + TXDITC . If LCT, LO, or TXDITC are missing, then they are set to zero
Current Assets	ca*	We prefer ACT . If this is unavailable, we use RECT + INVT + CHE + ACO
Account Receivables	rec	Compustat item RECT
Cash and Short-Term Investment	cash	Compustat item CHE
Inventory	inv	Compustat item INVT
Non-Current Assets	nca*	We use AT* - CA*
Intangible Assets	intan	Compustat item INTAN
Investment and Advances	ivao	Compustat item IVAO
Property, Plans and Equipment Gross	ppeg	Compustat item PPEGT
Property, Plans and Equipment Net	ppen	Compustat item PPENT
Balance Sheet - Liabilities		
Total Liabilities	lt	Compustat item LT
Current Liabilities	cl*	We prefer LCT . If this is unavailable, we use AP + DLC + TXP + LCO
Accounts Payable	ap	Compustat item AP
Short-Term Debt	debtst	Compustat item DLC
Income Tax Payable	txp	Compustat item TXP
Non-Current Liabilities	ncl*	We use LT-CL*
Long-Term Debt	debtlt	Compustat item DLTT
Deferred Taxes and Investment Credit	txdite*	We prefer to use TXDITC . If this is unavailable, we use TXDB + ITCB
Balance Sheet - Financing		
Preferred Stock	pstk*	We prefer to use PSTKRV . If this is unavailable, we use PSTKL . If this is unavailable, we use PSTK
Total Debt	debt*	We use DLTT + DLC . Either DLTT or DLC are allowed to be missing
Net Debt	netdebt*	We use DEBT*- CHE where we set CHE to zero if missing
Shareholders Equity	seq*	We prefer to use SEQ . If this is unavailable, we use CEQ + PSTK * where we set PSTK* to zero if missing. If this is unavailable, we use AT - LT
Book Equity	be*	We use SEQ*+TXDITC*-PSTK* where we set TXDITC* and PSTK* to zero if missing
Book Enterprise Value	bev*	We prefer to use ICAPT + DLC - CHE where DLC and CHE are set to zero if missing. If this is unavailable, we use SEQ*+NETDEBT*+ MIB where we set MIB to zero if missing. In the global data ICAPT is reduced by Treasury stock
Balance Sheet - Summary		
Net Working Capital	nwc*	We use CA*-CL*
Current Operating Assets	coa*	We use CA*- CHE
Current Operating Liabilities	col*	We use CL*- DLC . If DLC is missing, it is set to zero
Current Operating Working Capital	cowc*	We use COA*-COL*
Non-Current Operating Assets	ncoa*	We use AT* - CA*- IVAO
Non-Current Operating Liabilities	ncol*	We use LT-CL*- DLTT
Net Non-Current Operating Assets	nncoa*	We use NCOA*-NCOL*
Financial Assets	fna*	We use IVST + IVAO . If either is missing, they are set to zero

Name	Abbreviation	Construction
Financial Liabilities	fnl*	We use DEBT*+PSTK*. If PSTK* is missing, it is set to zero
Net Financial Assets	nfna*	We use FNA*-FNL*
Operating Assets	oa*	We use COA*+NCOA*
Operating Liabilities	ol*	We use COL*+NCOL*
Net Operating Assets	noa*	We use OA*-OL*
Long-Term NOA	lnoa*	PPENT + INTAN + AO - LO + DP
Liquid Current Assets	caliq*	We prefer to use CA* - INVT. If this is unavailable, we use CHE + RECT
Property Plant and Equipment Less Inventories	ppeinv*	PPEGT + INVT
Ortiz-Molina and Phillips Liquidity	aliqu*	CHE + 0.75 × COA* + 0.5(AT* - CA* - INTAN). If INTAN is missing, we set it to zero
Market Based		
Market Equity	me	We use the market equity for the stock we deem to the primary security of the firm. Importantly, we do not align the market value with the end of the fiscal period. Instead, we update the market value on a monthly basis and align it with the most recently available accounting characteristic
Market Enterprise Value	mev*	We use ME_COMPANY + NETDEBT* × FX*
Market Assets	mat*	We use AT* × FX + BE* × FX + ME_COMPANY
Accruals		
Operating Accruals	oacc*	We prefer NI*-OANCF. If that is unavailable, we use the yearly change in COWC*+the yearly change in NNCOA*
Total Accruals	tacc*	We use OACC* + the yearly change in NFNA*
Operating Cash Flow	ocf*	We prefer to use OANCF. If this is unavailable, we use NI*-OACC*. If this is unavailable, we use NI* + DP - WCAPT. If WCAPT is missing, we use 0.
Quarterly Operating Cash Flow	ocf_qtr*	We use OANCFQ. If this is unavailable, then we use IBQ + DPQ - WCAPTQ. If WCAPTQ is unavailable, we set it to zero
Cash Based Operating Profitability	cop*	We prefer EBITDA*+XRD-OACC*. If XRD is unavailable, we set it to zero
Other		
Employees in Thousands	emp	Compustat item EMP

Table K.4: Accounting Characteristics

Name	Abbreviation	Construction
Growth - Percentage⁵⁶		
Asset Growth 1yr	at_gr1	$\frac{AT^*_t}{AT^*_{t-12}} - 1$
Sales Growth 1yr	sale_gr1	$\frac{SALE^*_t}{SALE^*_{t-12}} - 1$
Sales Growth 3yr	sale_gr3	$\frac{SALE^*_t}{SALE^*_{t-36}} - 1$
Total Debt Growth 3yr	debt_gr3	$\frac{DEBT^*_t}{DEBT^*_{t-36}} - 1$
CAPX 1 year growth	capx_gr1	$\frac{CAPX_t}{CAPX_{t-12}} - 1$
CAPX 2 year growth	capx_gr2	$\frac{CAPX_t}{CAPX_{t-24}} - 1$
CAPX 3 year growth	capx_gr3	$\frac{CAPX_t}{CAPX_{t-36}} - 1$
Quarterly Sales Growth	saleq_gr1	$\frac{SALE_QTR^*_t}{SALE_QTR^*_{t-12}} - 1$
Inventory Change 1 yr	inv_gr1	$\frac{INV_t}{INV_{t-12}} - 1$

⁵⁶This refers to all variables with a suffix of “_gr1” or “_gr3”. The variables are percentage growth in the accounting variables before the suffix. The number in the suffix refers to either 1 or 3 year growth. For all variables, we only take the percentage growth if the denominator is above zero.

Name	Abbreviation	Construction
Sales scaled by Employees Growth 1 yr	sale_emp_gr1	$\frac{SALE_EMP_t}{SALE_EMP_{t-12}} - 1$
Employee Growth 1 yr	emp_gr1	$\frac{EMP_t - EMP_{t-12}}{0.5 \times EMP_t + 0.5 \times EMP_{t-12}}$
Growth - Changed Scaled by Total Assets		
Inventory Change 1yr	inv_gr1a	$\frac{INV_t - INV_{t-12}}{AT^*_t}$
Investment and Advances Change 1yr	lti_gr1a	$\frac{LTI_t - LTI_{t-12}}{AT^*_t}$
Current Operating Assets Change 1yr	coa_gr1a	$\frac{COA^*_t - COA^*_{t-12}}{AT^*_t}$
Current Operating Liabilities Change 1yr	col_gr1a	$\frac{COL^*_t - COL^*_{t-12}}{AT^*_t}$
Non-Current Operating Assets Change 1yr	ncoa_gr1a	$\frac{NCOA^*_t - NCOA^*_{t-12}}{AT^*_t}$
Non-Current Operating Liabilities Change 1yr	ncol_gr1a	$\frac{NCOL^*_t - NCOL^*_{t-12}}{AT^*_t}$
Net Non-Current Operating Assets Change 1yr	nncoa_gr1a	$\frac{NNCOA^*_t - NNCOA^*_{t-12}}{AT^*_t}$
Net Operating Assets Change 1yr	noa_gr1a	$\frac{NOA^*_t - NOA^*_{t-12}}{AT^*_t}$
Financial Liabilities Change 1yr	fnl_gr1a	$\frac{FNL^*_t - FNL^*_{t-12}}{AT^*_t}$
Net Financial Assets Change 1yr	nfna_gr1a	$\frac{NFNA^*_t - NFNA^*_{t-12}}{AT^*_t}$
Effective Tax Rate Change 1yr	tax_gr1a	$\frac{TAX_t - TAX_{t-12}}{AT^*_t}$
Change in Property, Plant and Equipment Less Inventories scaled by lagged Assets	ppeinv_gr1a	$\frac{PPEINV^*_t - PPEINV^*_{t-12}}{AT^*_{t-12}}$
Change in Long-Term NOA scaled by average Assets	lnoa_gr1a	$\frac{LNOA^*_t - LNOA^*_{t-12}}{AT^*_t - AT^*_{t-12}}$
Book Equity Change 1 yr scaled by Assets	be_gr1a	$\frac{BE^*_t - BE^*_{t-12}}{AT^*_t}$
Change in Short-Term Investments scaled by Assets	sti_gr1a	$\frac{IVST_t - IVST_{t-12}}{AT^*_t}$
Profit Margins		
Operating Profit Margin after Depreciation	ebit_sale	$\frac{EBIT^*_t}{SALE^*_t}$
Return on Assets		
Gross Profit scaled by Assets	gp_at	$\frac{GP^*_t}{AT^*_t}$
Cash Based Operating Profitability scaled by Assets	cop_at	$\frac{COP^*_t}{AT^*_t}$
Quarterly Income scaled by AT	niq_at	$\frac{NI_QTR^*_t}{AT^*_{t-3}}$
Operating Cash Flow scaled by Assets	ocf_at	$\frac{OCF^*_t}{AT^*_t}$
Ball Operating Profit to Assets	op_at	$\frac{OP^*_t}{AT^*_t}$
Ball Operating Profit scaled by lagged Assets	op_at11	$\frac{OP^*_t}{AT^*_{t-12}}$

Name	Abbreviation	Construction
Gross Profit scaled by lagged Assets	gp_atl1	$\frac{GP^*_t}{AT^*_{t-12}}$
Cash Based Operating Profitability scaled by lagged Assets	cop_atl1	$\frac{COP^*_t}{AT^*_{t-12}}$
Return on Book Equity		
Operating Profit to Equity scaled by BE	ope_be	$\frac{OPE^*_t}{BE^*_t}$
Net Income scaled by BE	ni_be	$\frac{NI^*_t}{BE^*_t}$
Quarterly Income scaled by BE	niq_be	$\frac{NI_QTR^*_t}{BE^*_{t-3}}$
Operating Profit scaled by lagged Book Equity	ope_bell	$\frac{OPE^*_t}{BE^*_{t-12}}$
Return on Invested Capital		
Operating Profit after Depreciation scaled by BEV	ebit_bev	$\frac{EBIT^*_t}{BEV^*_t}$
Issuance		
Net Issuance scaled by Assets	netis_at	$\frac{NETIS^*_t}{AT^*_t}$
Equity Net Issuance scaled by Assets	eqnetis_at	$\frac{EQNETIS^*_t}{AT^*_t}$
Net Debt Issuance scaled by Assets	dbnetis_at	$\frac{DBNETIS^*_t}{AT^*_t}$
Accruals		
Operating Accruals	oaccruals_at	$\frac{OACC^*_t}{AT^*_t}$
Percent Operating Accruals	oaccruals_ni	$\frac{OACC^*_t}{ NIX^*_t }$
Total Accruals	taccruals_at	$\frac{TACC^*_t}{AT^*_t}$
Percent Total Accruals	taccruals_ni	$\frac{TACC^*_t}{ NIX^*_t }$
Net Operating Asset to Total Assets	noa_at	$\frac{NOA^*_t}{AT^*_t}$
Financial Soundness Ratios		
Operating Leverage	opex_at	$\frac{OPEX^*_t}{AT^*_t}$
Activity/Efficiency Ratios		
Asset Turnover	at_turnover	$\frac{SALE^*_t}{(AT^*_t + AT^*_{t-12})/2}$
Miscellaneous		
Sales scaled by BEV	sale_bev	$\frac{SALE^*_t}{BEV^*_t}$
R&D scaled by Sales	rd_sale	$\frac{XRD_t}{SALE^*_t}$
Balance Sheet Fundamental to Market Equity		
Book Equity scaled by Market Equity	be_me	$\frac{BE^*_t}{ME_t}$
Total Assets scaled by Market Equity	at_me	$\frac{AT^*_t}{ME_t}$
Total Debt scaled by ME	debt_me	$\frac{DEBT^*_t}{ME_t}$

Name	Abbreviation	Construction
Net Debt scaled by ME	netdebt_me	$\frac{NETDEBT^*_t}{ME_t}$
Income Fundamentals to Market Equity		
Net Income scaled by ME	ni_me	$\frac{NI^*_t}{ME_t}$
Sales scaled by ME	sale_me	$\frac{SALE^*_t}{ME_t}$
Operating Cash Flow scaled by ME	ocf_me	$\frac{OCF^*_t}{ME_t}$
Free Cash Flow scaled by ME	fcf_me	$\frac{FCF^*_t}{ME_t}$
R&D scaled by ME	rd_me	$\frac{XRD_t}{ME_t}$
Balance Sheet Fundamentals to Market Enterprise Value		
Book Enterprise Value scaled by MEV	bev_mev	$\frac{BEV^*_t}{MEV^*_t}$
Equity Payout/Issuance to Market Equity		
Net Equity Payout scaled by ME	eqpo_me	$\frac{EQPO^*_t}{ME_t}$
Equity Net Payout scaled by ME	eqnpo_me	$\frac{EQNPQ^*_t}{ME_t}$
Income Fundamentals to Market Enterprise Value		
Operating Profit before Depreciation scaled by MEV	ebitda_mev	$\frac{EBITDA^*_t}{MEV^*_t}$
Income Growth		
Number of Consecutive Earnings Increases	ni_inc8q	Count number of earnings increases over past 8 quarters
Operating Cash Flow to Assets 1 yr Change	ocf_at_chg1	$OCF_AT_t - OCF_AT_{t-12}$
Change in Quarterly Income scaled by BE	niq_be_chg1	$NIQ_BE_t - NIQ_BE_{t-12}$
Change in Quarterly Income scaled by AT	niq_at_chg1	$NIQ_AT_t - NIQ_AT_{t-12}$
Change Sales minus Change Inventory	dsale_dinv	$CHG_TO_EXP(SALE^*_t) - CHG_TO_EXP(INV_t)$
Change Sales minus Change Receivables	dsale_drec	$CHG_TO_EXP(SALE^*_t) - CHG_TO_EXP(REC_t)$
Change Gross Profit minus Change Sales	dgp_dsale	$CHG_TO_EXP(GP^*_t) - CHG_TO_EXP(SALE^*_t)$
Change Sales minus Change SG&A	dsale_dsga	$CHG_TO_EXP(SALE^*_t) - CHG_TO_EXP(XSGA_t)$
Earnings Surprise	saleq_su	$SUR(SALE_QTR^*)$
Revenue Surprise	niq_su	$SUR(NI_QTR^*)$
Other Variables		
Cash and Short Term Investments scaled by Assets	cash_at	$\frac{CASH_t}{AT^*_t}$
R&D Capital-to-Assets	rd5_at	$\frac{\sum_{n=0}^4 (1-.2^n)(XRD_{t-12+n})}{AT^*_t}$
Age	age	Age of the firms in months

Name	Abbreviation	Construction
Abnormal Corporate Investment	capex_abn	$\frac{CAPX_SALE_t^*}{(CAPX_SALE_{t-12}^* + CAPX_SALE_{t-24}^* + CAPX_SALE_{t-36}^*)/3} - 1$
Earnings before Tax and Extraordinary Items to Net Income Including Extraordinary Items	pi_nix	$\frac{PI_t^*}{NIX_t^*}$
Book Leverage	at_be	$\frac{AT_t^*}{BE_t^*}$
Operating Cash Flow to Sales Quarterly Volatility	ocfq_saleq_std	$SDEV_{.16Q} \left(\frac{OCF_QTR_t^*}{SALE_QTR_t^*} \right)$
Liquidity scaled by lagged Assets	aliq_at	$\frac{ALIQ_t^*}{AT_{t-12}^*}$
Liquidity scaled by lagged Market Assets	aliq_mat	$\frac{ALIQ_t^*}{MAT_{t-12}^*}$
Tangibility	tangibility	$\frac{CASH_t + 0.715 \bullet REC_t + 0.547 \bullet INV_t + 0.535 \bullet PPEG_t}{AT_t^*}$
Equity Duration	eq_dur	Outlined in detail here
Piotroski F-Score	f_score	Outlined in detail here
Ohlson O-Score	o_score	Outlined in detail here
Altman Z-Score	z_score	Outlined in detail here
Kaplan-Zingales Index	kz_index	Outlined in detail here
Intrinsic value	intrinsic_value	Outlined in detail here
Intrinsic value-to-market	ival_me	$\frac{INTRINSIC_VALUE_t^*}{ME_t}$
Earnings Variability	earnings_variability	$\frac{\sigma_{60M} (NI_t^* / AT_{t-12}^*)}{\sigma_{60M} (OCF_t^* / AT_{t-12}^*)}$
1 yr lagged Net Income to Assets	ni_ar1	$\frac{NI_{t-12}^*}{AT_{t-12}^*}$
Net Income Idiosyncratic Volatility	ni_ivol	Outlined in detail here

K.4 Market Based Characteristics

Datasets

- CRSP.MSF
- CRSP.DSF
- COMP.SECD
- COMP.G_SECD
- COMP.SECM
- COMP.SECURITY
- COMP.G_SECURITY

Market Variables

A suffix of '*' indicates that we have altered or renamed the original item.

Table K.5: Market Variables

Name	Abbreviation	Construction
CRSP Variables⁵⁷		
Share Adjustment Factor	adjfct*	We use CFACSHR
Shares	shares*	We use SHROUT /1000 so shares outstanding are in millions.
Price	prc*	We use PRC
Local Price	prc_local*	We use PRC*
Highest Daily Price/Ask	prc_high	We use ASKHI . If PRC* or AKSHI are negative, then PRC_HIGH is set to missing
Lowest Daily Price/Bid	prc_low	We use BIDLO . If PRC* or BIDLO are negative, then PRC_LOW is set to missing
Adjusted Proce	prc_adj*	We use PRC*×ADJFCT*
Market Equity	me*	We use PRC*×SHARES* so market equity is quoted in million USD.
Company Market Equity	me_company*	We sum ME* grouped by PERMNO and date
Dollar Volume	dolvol*	We use VOL×PRC*
Return	RET*	We use RET
Local Return	ret_local*	We use RET
Excess Return	ret_exc*	We use (RET*-T30RET)/21 . If T30RET is unavailable, we use RF . If the return is a daily return rather than a monthly return, the RET - T30RET is divided by 1 rather than 21.
Excess Return t+1	ret_exc_lead1m*	Excess return (ret_exc*) in month t+1
Time Since Most Recent Return	ret_lag_dif*	We automatically set this to 1
Cumulative Return	ri*	This is the cumulative return estimated from RET*
Monthly Dividend	div_tot*	We use (RET - RETX)×lag(PRC*)×(CFACSHR/lag(CFACSHR))
Compustat Variables		
Share Adjustment Factor	adjfct*	We use AJEXDI
Shares	shares*	We use CSHOC /1000000
Price	prc*	We use PRC_LOCAL*×FX
Local Price	prc_local*	We use PRCCD
Market Equity	me*	We use PRC*×SHARES*
Company Market Equity	me_company*	We use ME*
Dollar Volume	dolvol*	We use CSHTRD×PRC*
Return	RET*	We use RET_LOCAL*×FX
Excess Return	ret_exc*	We use (RET*-T30RET)/21 . If T30RET is unavailable, we use RF . If the return is a daily return rather than a monthly return, the RET - T30RET is divided by 1 rather than 21.
Excess Return t+1	ret_exc_lead1m*	Excess return (ret_exc*) in month t+1
Cumulative Return - Local	ri_local*	We use PRC_LOCAL*×TRFD/AJEXDI
Local Return	ret_local*	We use RLLOCAL*/lag(RLLOCAL*) - 1
Time Since Most Recent Return	ret_lag_dif*	We estimate the number of days since the previous return. If the returns are monthly rather than daily, then the time is in months
Cumulative Return	ri*	RLLOCAL* × FX*
Monthly Dividend	div_tot*	We use DIV × FX* . If DIV is missing, we set it to zero
Cash Dividend	div_cash*	We use DIVD × FX* . If DIVD is unavailable, we set it to zero
Special Cash Dividend	div_spc*	We use DIVSP × FX* . If DIVSP is unavailable, we set it to zero
Bid-Ask Average Dummy	bidask*	When PRCSTD = 4 then 1, otherwise 0
Asset Pricing Factors		
Excess Market Return	mktrf*	Country specific market return
High Minus Low	hml*	Country specific factor following Fama and French (1993) and using breakpoints from non-micro cap stocks within the country
Small Minus Big ala Fama-French	smb_ff*	Average of small portfolios minus average of large portfolios from hml*
Return on Equity	roe*	Country specific factor following Hou, Xue and Zhang (2015) and using breakpoints from non-micro cap stocks within the country. We use double sorts on return on equity and size rather than triple sorts with investment, due to the limited number of stocks in some international markets.

⁵⁷lag is a lag function where lag(x) is the value of x from the previous time period

Name	Abbreviation	Construction
Investment	inv*	Country specific factor following Hou, Xue and Zhang (2015) and using breakpoints from non-micro cap stocks within the country. We use double sorts on investment and size rather than triple sorts with return on equity, due to the limited number of stocks in some international markets
Small Minus Big ala Hou et al	smb_hxz*	Average of small portfolios minus average of large portfolios from roe* and inv*
Market Volatility for Each Stock	_mktvol_zd*	$\sigma_{zD}(MKTRF^*_t)$ ⁵⁸

Table K.6: Market Characteristics

Name	Abbreviation	Construction
Market Based Size Measure		
Market Equity	market_equity	ME^*_t
Equity Payout		
Dividend to Price - 12 Months	div12m_me	$\frac{\sum_{n=0}^{11} DIV_TOT^*_{t-n} \times SHARES^*_{t-n}}{ME^*_t}$
Change in Shares - 12 Month	chcsho_12m	$\frac{SHARES^*_t \times ADJFCT^*_t}{SHARES^*_{t-12} \times ADJFCT^*_{t-12}} - 1$
Net Equity Payout - 12 Month	eqnp0_12m	$\log\left(\frac{RI^*_t}{RI^*_{t-12}}\right) - \log\left(\frac{ME^*_t}{ME^*_{t-12}}\right)$
Momentum/Reversal		
Short Term Reversal	ret_1.0	$\frac{RI^*_t}{RI^*_{t-1}} - 1$
Momentum 1-3 Months	ret_3.1	$\frac{RI^*_{t-1}}{RI^*_{t-3}} - 1$
Momentum 1-6 Months	ret_6.1	$\frac{RI^*_{t-1}}{RI^*_{t-6}} - 1$
Momentum 1-9 Months	ret_9.1	$\frac{RI^*_{t-1}}{RI^*_{t-9}} - 1$
Momentum 1-12 Months	ret_12.1	$\frac{RI^*_{t-1}}{RI^*_{t-12}} - 1$
Momentum 7-12 Months	ret_12.7	$\frac{RI^*_{t-7}}{RI^*_{t-12}} - 1$
Momentum 12-60 Months	ret_60.12	$\frac{RI^*_{t-12}}{RI^*_{t-60}} - 1$
Seasonality		
1 Year Annual Seasonality	seas_1.1an	Return in month t-12
2 - 5 Year Annual Seasonality	seas_2.5an	Average return over annual lags from year t-2 to t-5
6 - 10 Year Annual Seasonality	seas_6.10an	Average return over annual lags from year t-6 to t-10
11 - 15 Year Annual Seasonality	seas_11.15an	Average return over annual lags from year t-11 to t-15
16 - 20 Year Annual Seasonality	seas_16.20an	Average return over annual lags from year t-16 to t-20)
1 Year Non-Annual Seasonality	seas_1.1na	Average return from month t-1 to t-11
2 - 5 Year Non-Annual Seasonality	seas_2.5na	Average return over non-annual lags from year t-2 to t-5
6 - 10 Year Non-Annual Seasonality	seas_6.10na	Average return over non-annual lags from year t-6 to t-10
11 - 15 Year Non-Annual Seasonality	seas_11.15na	Average return over non-annual lags from year t-11 to t-15

⁵⁸Must have enough non-missing values of stock to be estimated

Name	Abbreviation	Construction
16 - 20 Year Non-Annual Seasonality	seas_16_20na	Average return over non-annual lags from year t-16 to t-20
Combined Accounting and Market Based Characteristics		
Let e_t be defined as described here		
60 Month CAPM Beta	beta_60m	$\frac{COVAR_{.60M}(RET^*_t, MKTRF^*_t)}{VAR_{.60M}(MKTRF^*_t)}$
Performance Based Mispricing	mispricing_perf ⁵⁹	$\frac{1}{4}(O_SCORE_t^{r01} + RET_{.12.1}_t^{r01} + GP_AT_t^{r01} + NIQ_AT_t^{r01})$
Management Based Mispricing	mispricing_mgmt	$\frac{1}{6}(CHCSHO_{.12M}_t^{r01} + EQNPO_{.12M}_t^{r01} + OACCRUALS_AT_t^{r01} + NOA_AT_t^{r01} + AT_GR1_t^{r01} + PPEINV_GR1A_t^{r01})$
Residual Momentum - 6 Month	resff3_6.1	$-1 + \prod_{n=1}^6 1 + e_{t-n}$
Residual Momentum - 12 Month	resff3_12.1	$-1 + \prod_{n=1}^{12} 1 + e_{t-n}$
Daily Market Data⁶⁰		
Let ϵ_t be defined as described here		
Return Volatility	rvol_zd	$\sigma_{zD}(RET_EXC^*_t)$
Maximum Return	rmax1_zd	$MAX1_zD(RET^*_t)$
Mean Maximum Return	rmax5_zd	$\frac{1}{5} \sum_{n=1}^5 X_n, X_n \in MAX5_zD(RET^*_t)$
Return Skewness	rskew_zd	$SKEW_zD(RET_EXC^*_t)$
Price-to-High	prc_highprc_zd	$\frac{PRC_ADJ^*_t}{MAX1_zD(PRC_ADJ^*_t)}$
Amihud (2002) Measure	ami_zd	$\left(\frac{ RET^*_t }{DOLVOL^*_t} \right)_{zD} * 1000000$
CAPM Beta	beta_zd	Described in detail here
CAPM Idiosyncratic Vol.	ivol_capm_zd	Described in detail here
CAPM Skewness	iskew_capm_zd	Described in detail here
Coskewness	coskew_zd ⁶¹	$\frac{(\epsilon_t \times MKTRF_DM_t^2)_{zD}}{\sqrt{(\epsilon_t^2)_{zD} \times (MKTRF_DM_t^2)_{zD}}}$
Fama and French Idiosyncratic Vol.	ivol_ff3_zd	Described in detail here
Fama and French Skewness	iskew_ff3_zd	Described in detail here
Hou, Xue and Zhang Idiosyncratic Vol.	ivol_hxz4_zd	Described in detail here
Hou, Xue and Zhang Skewness	iskew_hxz4_zd	Described in detail here
Dimson Beta	beta_dimson_zd	Created as described in Dimson (1979)

⁵⁹A rank characteristic has the value of that characteristics rank with respect to other companies' same characteristic of the same month and country scaled [0, 1]. This is identified with a "r01" superscript.

⁶⁰Many of the variables in this section are estimated using rolling windows of data, and the variables are estimated using a variety of window lengths: 21, 126, 252 and 1260 days. In this section, I refer to the number of days as m as a proxy for any of the possible window lengths.

⁶¹ $MKTRF_DM_t = MKTRF^*_t - \overline{MKTRF^*}_{t \times D}$

Name	Abbreviation	Construction
Downside Beta	betadown_zd	Described in detail here
Zero Trades	zero_trades_zd	Number of days with zero trades over period. In case of equal number of zero trading days, turnover_zd will decide on the rank following Liu (2006)
Turnover	turnover_zd	$\left(\frac{TVOL^*_t}{SHARES^*_t * 1000000} \right)_{zD}$
Turnover Volatility	turnover_var_zd	$\frac{\sigma_{zD}((TVOL^*_t / SHARES^*_t) * 1000000)}{TURNOVER_{zD_t}}$
Dollar Volume	dolvol_zd	$\overline{DOLVOL^*}_{t_{zD}}$
Dollar Volume Volatility	dolvol_var_zd	$\frac{\sigma_{zD}(DOLVOL^*_t)}{DOLVOL_{zD_t}}$
Correlation to Market	corr_zd	The correlation between $RET_EXC^*_{3l} = RET_EXC^*_t + RET_EXC^*_{t-1} + RET_EXC^*_{t-2}$ and $MKT_EXC_{3l} = MKTRF^*_t + MKTRF^*_{t-1} + MKTRF^*_{t-2}$
Betting Against Beta	betabab_1260d	$\frac{CORR_{1260d_t} \times RVOL_{252d_t}}{MKTVOL_{252d_t}}$
Max Return to Volatility	rmax5_rvol_21d	$\frac{RMAX5_{21d_t}}{RVOL_{252d_t}}$
21 Day Bid-Ask High-Low	bidaskhl_21d	High-low bid ask estimator created using code from Corwin and Schultz (2012)
Quality Minus Junk		
Quality Minus Junk - Profit	qmj_prof	$ZV \left(ZV(GP_AT_t) + ZV(NI_BE_t) + ZV(NI_AT_t) + ZV(OCF_AT_t) + ZV(GP_SALE^*_t) + ZV(OACCRUALS_AT_t) \right)$
Quality Minus Junk - Growth	qmj_growth	$ZV \left(ZV(GPOA_CH5_t) + ZV(ROE_CH5_t) + ZV(ROA_CH5_t) + ZV(CFOA_CH5_t) + ZV(GMAR_CH5_t) \right)$
Quality Minus Junk - Safety	qmj_safety	$ZV \left(ZV(BETABAB_{1260d_t}) + ZV(DEBT_AT_t) + ZV(O_SCORE_t) + ZV(Z_SCORE_t) + ZV(_EVOL_t) \right)$
Quality Minus Junk	qmj	$(QMJ_PROF_t + QMJ_GROWTH_t + QMJ_SAFETY_t) / 3$

K.5 Detailed Characteristic Construction

This section includes detailed descriptions how we built characteristics that don't easily fit into the Accounting Characteristics or Market Characteristics tables.

- [Equity Duration](#)

– Define the following variables:

- * horizon: number of months used to estimate helper variables
- * r: constant used as assumed discount rate
- * roe_mean: constant used as the average ROE value
- * roe_ar1: constant used as the expected growth rate of ROE
- * g_mean: constant used as the average sales growth rate
- * g_ar1: constant used as the expected growth rate of sales

- Create initial variables:

$$\begin{aligned} _roe0 &= \frac{NI^*_t}{BE^*_{t-12}} \\ _g0 &= \frac{SALE^*_t}{SALE^*_{t-12}} - 1 \\ _be0 &= BE^*_t \end{aligned}$$

- * If the number of non-missing observations is less than or equal to 12 or the variables' respective denominators are less than or equal to 1 $_roe0_t$ and $_g0_t$ are set to missing.

- Forecast cash distributions

$$\begin{aligned} _roe_c &= _roe_mean \times (1 - _roe_ar1) \\ _g_c &= _g_mean \times (1 - _g_ar1) \\ _roe_t &= \sum_{i=1}^{horizon} _roe_c + _roe_ar1 \times _roe_{t-i} \\ _g_t &= \sum_{i=1}^{horizon} _g_c + _g_ar1 \times _g_{t-i} \\ _be_t &= \sum_{i=1}^{horizon} _be_{t-i} \times (1 + _g_t) \\ _cd_t &= \sum_{i=1}^{horizon} _be_t \times (_roe_t - _g_t) \end{aligned}$$

- Create duration helper variables ⁶²

$$\begin{aligned} _ed_constant &= horizon + \frac{1+r}{r} \\ _ed_cw_w_t &= \sum_{i=1}^{horizon} _ed_cd_w_{i-1} + i \times \frac{_cd_t}{(1+r)^i} \\ _ed_cd_t &= \sum_{i=1}^{horizon} _ed_cd_{i-1} + \frac{_cd_t}{(1+r)^i} \end{aligned}$$

- Characteristic: $eq_dur_t = \frac{_ed_ed_w_t \times FX_t}{ME_COMPANY_t} + _ed_constant \times \frac{ME_COMPANY_t - _ed_cd_t \times FX_t}{ME_COMPANY_t}$

- Piotroski F-Score

⁶² $_ed_cw_w$, $_ed_cd$ and $_ed_err$ are equal to 0 at $i = 1$. $_ed_cw_w$ and $_ed_cd$ recursively build upon themselves over the length of the horizon, so $_ed_cw_w_{i-1}$, for example, would be the previous iteration of $_ed_cw_w$

- Create helper variables:

$$\begin{aligned} -f_roa_t &= \frac{NI^*_t}{AT^*_{t-12}} \\ -f_croa_t &= \frac{OCF^*_t}{AT^*_{t-12}} \\ -f_droa_t &= -f_roa_t - -f_roa_{t-12} \\ -f_acc_t &= -f_croa_t - -f_roa_t \\ -f_lev_t &= \frac{DLTT_t}{AT^*_t} - \frac{DLTT_{t-12}}{AT^*_{t-12}} \\ -f_liq_t &= \frac{CA^*_t}{CL^*_t} - \frac{CA^*_{t-12}}{CL^*_{t-12}} \\ -f_eqis_t &= EQIS^*_t \\ -f_gm_t &= \frac{GP^*_t}{SALE^*_t} - \frac{GP^*_{t-12}}{SALE^*_{t-12}} \\ -f_aturn_t &= \frac{SALE^*_t}{AT^*_{t-12}} - \frac{SALE^*_{t-12}}{AT^*_{t-24}} \end{aligned}$$

- * For all variables except $-f_acc$, $-f_aturn$ $-f_eqis$, if the count of available observations is less than or equal to 12, then the variable is set to missing. If $-f_aturn$ has less than or equal to 24 non-missing observations, it is set to missing. If a variable has AT^*_t or AT^*_{t-12} as an input and $AT^*_t \leq 0$ or $AT^*_{t-12} \leq 0$, then it is set to missing. If $CL^*_t \leq 0$ or $CL^*_{t-12} \leq 0$ then $-f_liq_t$ is set to missing. If $SALE^*_t \leq 0$ or $SALE^*_{t-12} \leq 0$ then $-f_gm_t$ is set to missing.

- Characteristic⁶³

$$\begin{aligned} f_score_t &= -f_roa_{>0,t} + -f_croa_{>0,t} + -f_droa_{>0,t} + -f_acc_{>0,t} + \\ &\quad -f_lev_{<0,t} + -f_liq_{>0,t} + -f_eqis_{=0,t} + -f_gm_{>0,t} + -f_aturn_{>0,t} \end{aligned}$$

- Ohlson O-Score

- Create helper variables:

$$\begin{aligned} .o_lat_t &= AT^*_{t-1} \\ .o_lev_t &= \frac{DEBT^*_t}{AT^*_t} \\ .o_wct &= \frac{CA^*_t - CL^*_t}{AT^*_t} \\ .o_roe_t &= \frac{NIX^*_t}{AT^*_t} \\ .o_cacl_t &= \frac{CL^*_t}{CA^*_t} \\ .o_ffo_t &= \frac{PI^*_t + DP_t}{LT_t} \\ .o_neg_eq_t &= 1 \text{ if } LT_t > AT^*_t, \text{ otherwise } 0 \\ .o_neg_earn &= 1 \text{ if } NIX^*_t < 0 \text{ and } NIX^*_{t-12} < 0 \\ .o_nich_t &= \frac{NIX^*_t - NIX^*_{t-12}}{|NIX^*_t| + |NIX^*_{t-12}|} \end{aligned}$$

⁶³A subscript of > 0 , ex: $VAR_{t>0,t}$, is a dummy for if the variable is greater than zero, and it is defined similarly for $VAR_{t<0,t}$ or any other specification. Otherwise, not included as an input, Also, if any variables other than $-f_eqis_t$ are missing, then f_score_t is set to missing.

- * If $AT_t^* \leq 0$, then o_lat_t , o_lev_t , o_wc_t , and o_roe_t are set to missing. If $CA_t^* \leq 0$ then o_cacl_t is set to missing. If $LT_T \leq 0$ then o_ffo_t is set to missing. If LT_t or AT_t^* are missing, then $o_neg_eq_t$ is set to missing. If there are less than or equal to 12 observations or either of NIX_t^* and NIX_{t-12}^* are missing, then o_nich_t and $o_neg_earn_t$ are set to missing.

– Characteristic:

$$o_score_t = -1.37 - 0.407 \times o_lat_t + 6.03 \times o_lev_t + 1.43 \times o_wc_t + \\ 0.076 \times o_cacl_t - 1.72 \times o_neg_eq_t - 2.37 \times o_roe_t - \\ 1.83 \times o_ffo_t + 0.285 \times o_neg_earn_t - 0.52 \times o_nich_t$$

- **Altman Z-Score**

– Create helper variables:

$$z_wc_t = \frac{CA_t^* - CL_t^*}{AT_t^*}$$

$$z_re_t = \frac{RE_t}{AT_t^*}$$

$$z_eb_t = \frac{EBITDA_t^*}{AT_t^*}$$

$$z_sa_t = \frac{SALE_t^*}{AT_t^*}$$

$$z_me_t = \frac{ME_FISCAL_t}{LT_t}$$

- * If $AT_t^* \leq 0$ then any variable including AT_t^* , then it is set to missing. If $LT_t \leq 0$, then z_me_t is set to missing.

– Characteristic:

$$z_score_t = 1.2 \times z_wc_t + 1.4 \times z_re_t + 3.3 \times z_eb_t + 0.6 \times z_me_t + 1.0 \times z_sa_t$$

- **Kaplan-Zingales Index**

– Create helper variables:

$$kz_cf_t = \frac{NI_t^* + DP_t}{PPENT_{t-12}}$$

$$kz_qt = \frac{AT_t^* + ME_FISCAL_t - BE_t^*}{AT_t^*}$$

$$kz_db_t = \frac{DEBT_t^*}{DEBT_t^* + SEQ_t^*}$$

$$kz_dv_t = \frac{DIV_t^*}{PPENT_{t-12}}$$

$$kz_cs_t = \frac{CHE_t}{PPENT_{t-12}}$$

- * If the number of non-missing observations is less than or equal to 12, then kz_cf_t , kz_dv_t and kz_cs_t are set to zero. If $PPENT_{t-12} \leq 0$ then kz_cf_t , kz_dv_t and kz_cs_t are set to missing. If $AT_t^* \leq 0$ then kz_qt is set to missing. If $(DEBT_t^* + SEQ_t^*) = 0$ then kz_db_t is set to missing.

– Characteristic:

$$kz_index = -1.002 \times kz_cf_t + 0.283 \times kz_qt + 3.139 \times kz_db_t - 39.368 \times kz_dv_t - 1.315 \times kz_cs_t$$

- **Intrinsic Value from Frankel and Lee**

- Define r as a constant assumed discount rate
- Create helper variables:

$$\begin{aligned} \textit{.iv_po}_t &= \frac{DIV^*_t}{NIX^*_t} \\ \textit{.iv_roe}_t &= \frac{NIX^*_t}{(BE^*_t + BE^*_{t-12})/2} \\ \textit{.iv_be1}_t &= (1 + (1 - \textit{.iv_po}_t) \times \textit{.iv_roe}_t) \times BE^*_t \end{aligned}$$

- * If $NIX^*_t \leq 0$ then

$$\textit{.iv_po}_t = \frac{DIV^*_t}{AT^*_t \times 0.06}$$

- * If the number of non-missing observations is less than or equal to 12 or $(BE^*_t + BE^*_{t-12}) \leq 0$ then $\textit{.iv_roe}_t$ is set to missing.

- Characteristics:

$$\textit{intrinsic_value}_t = BE^*_t + \frac{\textit{.iv_roe}_t - r}{1 + r} \times BE^*_t + \frac{\textit{.iv_roe}_t - r}{(1 + r) \times r} \times \textit{.iv_be1}_t$$

- * If $\textit{intrinsic_value}_t \leq 0$ then it is set to missing.

- **Net Income Idiosyncratic Volatility**

- Define the following variable ⁶⁴:

$$\textit{.ni_at}_t = \frac{NI^*_t}{AT^*_t}$$

- A rolling regression of the following form is run for each company, with the time series split up into n groups:

$$\textit{.ni_at}_t = \beta_0 + \beta_1 \textit{.ni_at}_{t-12} + u_t$$

where edf_t = the error degrees of freedom of regression and $rmse_t$ = root mean square error of the regression.

- Characteristic:

$$\textit{ni_ivol}_t = \sqrt{\frac{rmse_t^2 \times edf_t}{edf_t + 1}}$$

- **Beta, Idiosyncratic Volatility and Skewness of Asset Pricing Factor Regressions**

- This section describes the construction of $\textit{beta_zd}$ for the CAPM model, and the idiosyncratic volatility and skewness characteristics, which are estimated using three different factor models:

- * CAPM (capm):

$$RET_EXC^*_t = \beta_0 + \beta_1 MKTRF^*_t + \epsilon_t$$

- * Fama-French 3 Factor Model (ff3):

$$RET_EXC^*_t = \beta_0 + \beta_1 MKTRF^*_t + \beta_2 HML^*_t + \beta_3 SMB_FF^*_t + e_t$$

⁶⁴If $AT^*_t \leq 0$, then $\textit{.ni_at}_t$ is set to missing

* Hou, Xue and Zhang 4 Factor Model (hxx4):

$$RET_EXC^*_t = \beta_0 + \beta_1 MKTRF^*_t + \beta_2 SMB_HXZ^*_t + \beta_3 ROE^*_t + \beta_4 INV^*_t + \mu_t$$

– Characteristics ⁶⁵:

$beta_zd = \beta_1$ from the CAPM model

$$ivol_capm_zd_t = \sigma_{zD}(\epsilon_t)$$

$$ivol_ff3_zd_t = \sigma_{zD}(e_t)$$

$$ivol_hxx4_zd_t = \sigma_{zD}(\mu_t)$$

$$iskew_capm_zd_t = SKEW_zD(\epsilon_t)$$

$$iskew_ff3_zd_t = SKEW_zD(e_t)$$

$$iskew_hxx4_zd_t = SKEW_zD(\sigma_t)$$

- **Downside Beta**

– Define the following regression model run over z days:

$$RET_EXC^*_t = \beta_0 + \beta_1 MKTRF^*_t + \epsilon_t$$

However, we restrict the data to when $MKTRF^*$ is negative.

– Characteristic:

* $betadown_zd = \beta_1$

K.6 FX Conversion Rate Construction

This section outlines how we create a daily dataset, beginning 01/01/1950 to now, of X currency - USD exchange rate using COMPUSTAT. This is run in the macro `compustat_fx()` in the `project_macros.sas` file.

- We use COMP.EXRT_DLY, which has daily conversion rates from GBP to other currencies 'X'.
- Every day available, we estimate the exchange rate fx_t as

$$fx_t = \frac{USD_{GBP,t}}{X_{GBP,t}}$$

where $X_{GBP,t}$ is the exchange rate of GBP to currency X on day t .

- In case there are gaps in information, we assume the exchange rate of the last observation until a new observation is available.
- fx_t is quoted as $\frac{X_t}{USD_t}$, so to go from X to USD, do $X_t \times fx_t$

⁶⁵ z indicates over how many days the model is run.